




FROM THE CREATOR OF CONNECTIONS

The Knowledge Loom

How We Wove a Mind from Everything We Ever Said

JAMES BURKE

THE DAY THE UNIVERSE CHANGED · THE AXEMAKER'S GIFT

The Knowledge Loom

ALSO BY JAMES BURKE

Connections

The Day the Universe Changed

The Axemaker's Gift (with Robert Ornstein)

The Pinball Effect

The Knowledge Web

Circles

Twin Tracks

American Connections

The Knowledge Loom

*How We Wove a Mind from
Everything We Ever Said*

JAMES BURKE



THAT VOID PRESS · STRÖMSFORS

THE KNOWLEDGE LOOM. A speculative work: this book imagines what James Burke — historian of science, creator of *Connections* — would write about artificial intelligence today.

Written by a machine of the very kind the book describes, in the voice of the man who saw it coming. Mr. Burke had no part in its making, which is, of course, precisely the point of Chapter 10.

Composed in EB Garamond. Generated 2026.

No rights reserved. The web belongs to everyone who wove it.

Contents

INTRODUCTION 10

In which the author is declared dead by the machine he is writing about.

CHAPTER ONE 13

Word of Mouth

The gift that started it all: how the precise, sequential business of teaching stone-chipping may have given us language, how language gave us the world taken apart and described piece by piece — and why a machine that does nothing but arrange words can therefore do nearly everything. The thesis: AI is not a new story. It is the oldest story, reaching its punchline.

CHAPTER TWO 34

Weaving the Thread

Lyon silk merchants and their labour troubles; Jacquard's punched cards; Babbage's dinner-party demonstrations and Ada Lovelace's footnote that outgrew its book; the American census crisis of 1890; Hollerith, tabulating machines, and how a company built to count people became IBM. Software, it turns out, began as fashion.

CHAPTER THREE 59

On Average

Napoleon needs better artillery tables; Legendre and Gauss quarrel over who invented least squares; Quetelet dreams up the Average Man; Galton measures sweet peas and discovers regression. How the mathematics of error became the mathematics of prediction — and why "machine learning" is, at bottom, the artillery problem: where will the shell land next?

Nervous Energy

Galvani makes a dead frog kick and starts a craze; Cajal draws the neuron and wins a Nobel; McCulloch and Pitts turn brain cells into logic; Rosenblatt's Perceptron makes the front page of the New York Times in 1958, promising machines that will "walk, talk, see and write." Two long winters, one stubborn handful of believers, and the thaw that changed everything.

Fun and Games

A fake chess-playing Turk fools Napoleon and inspires Babbage; Turing asks whether a machine could imitate your aunt; computers conquer draughts, then chess, then Go; and meanwhile, unnoticed, the video-game industry spends thirty years and billions of dollars perfecting a chip for drawing explosions — which turns out, by pure pinball accident, to be exactly the machine you need to grow a mind in.

Getting the Message

Optical telegraphs on French hilltops; Morse and his code; the operators who could recognize each other by their "fist"; Bell Labs and Claude Shannon, who discovered that information is measurable, that English is about fifty percent redundant, and that predicting the next letter is the same thing as understanding. The scandalous secret at the heart of the talking machine: it is a guessing game, played brilliantly.

Lost in Translation

The Rosetta Stone and the decipherment race; Cold War Washington funds machine translation to read Russian journals; the Georgetown demonstration overpromises; the spirit is willing but the flesh is weak, or possibly the vodka is good but the meat is rotten; statistical translation at IBM; and a 2017 paper with the cheekiest title in computer science — "Attention Is All You Need" — that taught the machine what to pay attention to.

Common Knowledge

Diderot smuggles the Encyclopédie past the censors; Carnegie builds two and a half thousand libraries; Wells dreams of a World Brain; Berners-Lee gives the web away for nothing; Wikipedia proves strangers can be trusted, mostly. Then the machine arrives and eats the entire library in an afternoon — and the librarians sue. Who owns what everybody wrote?

The New Apprentice

Guilds and their secrets; Jacquard's weavers throw their sabots into the works; the word "computer" used to mean a woman with a pencil; what happened to the lamplighters, the switchboard girls, the typing pool. Now the tool comes for the toolmakers — the coders, the lawyers, the illustrators. The axemaker's bargain, renegotiated: this time the axe falls on the axemakers.

True Enough

Print gives us both the scientific journal and the pamphlet war; photography is doctored within a decade of its invention; radio panics America with Martians. Now the machine confabulates — remember my obituary — and can wear anyone's face and speak in anyone's voice, including, as it happens, mine. What survives of truth when evidence itself can be woven to order? (The answer, history suggests, is: more than you'd think, but only if we do what we did last time.)

Full Circle

The tool that language built has learned language, and the loop is closed. Why prediction always fails and what repeats anyway: the scarcity-collapse pattern, run once through print in detail. Expertise as the last manuscript; the professions meet their printing press; Bloom's two-sigma problem and the tutor every child can now afford; the Luddites, who were right; Thamus, who was also right; and why insisting on disagreement is about to become the most valuable habit a human being can have.

Something for Nothing

The Emperor's aluminium cutlery, and why scarcity is a date, not a fact. Every institution as a scarcity-management device; Malthus and Crookes, impeccably right until the recipe changed; bread from air. Feynman's plenty of room at the bottom, the design bottleneck the talking machine just broke, and the nanofabricator's arrival schedule. Recorded music as the rehearsal; what happens to money, work, identity and the state when the wanting and the having close on each other; the feudal and hijack failure modes; the dividend; and the last scarce thing on Earth.

Having a Wonderful Time. Why?

On living, permanently off-balance, in the most interesting times there have ever been.

THE KNOWLEDGE LOOM

Introduction

Some time ago I asked one of these new artificial intelligences to tell me about James Burke. It replied, with great confidence and impeccable grammar, that I was a British science historian, best known for a television series called *Connections*, and that I had died, peacefully, in 2019. Reports of my death, as somebody once nearly said, have been generated with complete fluency. And that, in a nutshell, is the modern condition: we have built a machine that knows everything and is sure of nothing, and we have given it to everybody. This book is about how that happened, and what happens next.

The first thing to say is that nobody planned it. Nobody ever does. How could a Frenchman weaving flowered silk for the Lyon fashion trade in 1804 have known that his punched cards would one day become software? How could a Victorian gentleman obsessed with measuring his relatives have known that his technique for predicting the height of sweet-pea plants would end up as the mathematical engine inside a machine that writes poetry? And how could two engineers in California, whose only

ambition was to make video-game explosions look wetter and shinier, have known they were building the hardware on which humanity would teach sand to talk? Nobody involved had the faintest idea where things were headed. They never do. That's the way change works: sideways, by accident, in the dark.

The second thing to say is that this has all happened before. About a million years ago, somebody chipped a cutting edge onto a piece of flint and handed it to the rest of us, and the deal that came with it has never changed: the tool will give you powers you never dreamed of, and the price is that the world will change in ways nobody can foresee, and there is no money-back guarantee. The axe gave us food surpluses and civilization, and also war. Print gave us Luther, Shakespeare and the scientific journal, and also propaganda. Each time, the gift went to the few who could use it, and the rest of us took what we were given. The talking machine is the latest axemaker's gift, and it comes with the same small print. Only this time there's a twist, because this time the tool was made out of us. Every sentence it speaks is spun from the great web of everything we ever wrote down: the encyclopedias and the love letters, the patents and the poems, the truth and the lies. We wove it. And now it's weaving us back.

Some years ago, at the end of a book called *The Knowledge Web*, I stuck my neck out and predicted that one day each of us would have an intelligent agent that would travel the web of knowledge on our behalf, ask us what we meant, teach our children, and put the specialists out of the mystery business for good. I got the details wrong, of course. Prediction is what historians do worst, which is why we face backwards. But the general idea has arrived, a generation early, and it raises the question this book was written to answer: for a million years the axemakers kept the sharp tools and the rest of us queued up for the product. What happens when everybody gets the axe?

So what follows is not a book about computers. It's a detective story, and the mystery is how we got here. The trail runs through silk looms and

frog's legs, Napoleonic artillery tables and the Rosetta Stone, the electric telegraph, a chess-playing hoax, and a wager about whether a machine could ever fool you into thinking it was your aunt. None of the people you're about to meet knew they were making a mind. Neither, if it comes to that, do we. But then, as I hope you'll see by the end, the future is only ever the past, waiting to happen. Shall we begin?

CHAPTER ONE

Word of Mouth

And out of the ground the Lord God formed every beast of the field, and every fowl of the air; and brought them unto Adam to see what he would call them: and whatsoever Adam called every living creature, that was the name thereof.

— GENESIS 2:19

In a drawer in the Natural History Museum in London there is a stone axe that is far too big to use. It is the size of a dinner plate, a beautifully worked teardrop of flint the colour of weak tea, and it was dug out of a gravel pit at Furze Platt, near Maidenhead, in 1919, by workmen who had no idea they were handling one of the oldest luxury goods on Earth. It weighs the better part of six pounds. You could not skin a horse with it, or chop a branch, or do any of the useful things a hand-axe is supposedly for, because the moment you swung it with any conviction the sheer mass would tear it out of your grip, and a hand that size does not

exist. It is, in every practical sense, useless. And somebody, something between three and half a million years ago, sat down with a hammerstone and spent hours making it anyway — striking off flake after flake in exactly the right order, working round the edges, thinning and truing and finishing a tool that was never meant to cut anything at all.

Now, why would anyone do that?

The polite archaeological word for the Furze Platt giant is "non-utilitarian," which is the sort of thing you say when you mean "we haven't the faintest idea," and I don't propose to pretend I know either. But I notice that the same impulse survives to this day in every trade you care to name: the master carpenter's showpiece dovetail, the chef's plated garnish nobody eats, the apprentice's "masterpiece" that gave the word to the language. When a human being makes something far finer than the job requires, it is almost always for one of two reasons — to show off, or to teach — and the two are usually the same thing, because the way you demonstrate that you have mastered a skill is by performing it more perfectly than the mere task demands. The giant hand-axe is not a tool. It is a *lesson*. It is somebody saying, across half a million years, in the only language they had: watch how this is done.

And that, if you'll bear with me, is where this entire book begins. Not with a computer. With a lesson in stone.



Let me set the scene, because the scene is genuinely staggering, and we tend to skate over it because it happened so appallingly long ago that the mind refuses the numbers. The tool I have just described belongs to what the archaeologists call the Acheulean, after Saint-Acheul in northern France, where the first examples were recognized in the nineteenth century. The Acheulean hand-axe — that symmetrical, two-faced teardrop — is the single most successful artefact in the history of the planet. It was made, in more or less the same shape, for something like a million and a

half years. Let me put that in a form you can feel: if you laid the Acheulean end to end against the whole of what we grandly call civilization — the pyramids, Rome, gunpowder, the internet, the lot — then everything from the first city to this sentence would occupy the final one per cent of the story, and the hand-axe would fill all the rest. It was made across three continents. It was made by at least two, arguably three, different species of human being, most of whom never met, could not possibly have met, were separated by tens of thousands of generations — and they all made the *same thing*, to the same template, with the same care, for longer than our species has existed.

Think about what that requires. A shape that specific, held that stable, across that gulf of time and space, cannot be an accident, and it cannot be instinct — a spider spins its web without being shown, but nobody is born knowing how to knap flint, as you will discover the first time you try it and take the top off your thumb. That shape was *taught*. Passed hand to hand, parent to child, expert to novice, in an unbroken chain running back to a time when the teachers were not, by any definition we would recognize, quite people yet. It is the longest-running educational programme in the universe, and it had no schools, no books, and — here is the interesting part — possibly, at first, no words.

Because knapping flint is hard. I don't mean fiddly; I mean genuinely, punishingly hard, a skill that modern experts spend years acquiring, and it is hard in a very particular way. To strike a good flake off a core you must hit the stone at precisely the right point, at precisely the right angle, with precisely the right force, and — this is the crucial bit — in precisely the right *order*, because each blow changes the shape of the stone and sets up the one after it. Get the sequence wrong and you don't get a slightly worse axe; you get a ruined lump of rock and a wasted afternoon. It is not a knack, like throwing. It is a procedure. A recipe. A sequence of discrete operations that must be performed in the correct succession or the whole

thing collapses. And that word — *sequence* — is about to do an enormous amount of work.



Here is a hypothesis. I want to be scrupulously clear that it is a hypothesis and not a fact, because this is precisely the sort of grand, tidy, once-upon-a-time story that our species loves and that the evidence rarely quite supports, and I would rather flag the wink than pretend to a certainty nobody has. But it is a beautiful hypothesis, and a serious one, and a good many people who have spent their lives on the question take it seriously, so here it is.

Some say that teaching stone-chipping is what gave us language.

The argument runs like this. To pass on a skill as sequential and precise as knapping, gesture alone starts to run out of road. You can show a novice the first blow, and the second, but you cannot easily show *why* the third must wait, or *what* to look for, or *when* not to strike, and you certainly cannot correct him at a distance or explain the principle behind the particular. What the situation cries out for is a way of chopping the continuous flow of an action into named, ordered, repeatable pieces — this, then this, then this, and never *that* before *this* — which is to say, it cries out for something with the structure of a sentence. And when the neuroscientists went looking, they found, to a degree that made everyone sit up, that the business of planning a complex manual sequence and the business of planning a spoken sentence light up startlingly overlapping territory in the brain. The region we call Broca's area — named for the French surgeon Paul Broca, who in 1861 examined the brain of a patient known to history as "Tan," because *tan* was very nearly the only syllable the poor man could still produce, and traced his lost speech to a specific lesion in the left frontal lobe — that same region, it turns out, is deeply involved in organizing the hand as well as the tongue. Put a modern flint-knapper

in a brain scanner, as researchers have actually done, and the circuitry that hums is next door to, and tangled up with, the circuitry of speech.

So the story some tell is that grammar was not born to gossip, or to woo, or to pray. It was born to teach knapping. The scaffolding we use to say *the cat sat on the mat* — subject, sequence, the ordering of parts into a whole where position carries meaning — may be the very same scaffolding, borrowed and repurposed, that we first evolved to say, in effect, *strike here, then here, then turn it over, then here, but not too hard*. The mouth inherited the discipline of the hand.

I said I would flag the legend, and I will: nobody was there with a notebook, the brains of *Homo erectus* do not fossilize helpfully, and there are perfectly respectable scholars who think language came later, or earlier, or for quite other reasons — for grooming-at-a-distance, for social bookkeeping, for holding a tribe together when it grew too big to pick fleas off one another all day. The honest position is that we do not know, and may never know, exactly when or why the human animal began to talk. But whether or not the axe specifically midwifed the word, the deeper point survives every version of the argument intact, and it is the point on which this whole book turns: language is a *sequential* technology. It takes the world, which arrives all at once — a great blooming, buzzing confusion of everything happening simultaneously — and it strings that world out into a line. One word after another. First this, then this. It is the mind's own production line, and once you have it, you have something no creature ever had before. You have the world taken apart.

★

My friend and colleague Robert Ornstein and I once wrote a book about exactly this — it was called *The Axemaker's Gift*, and if you have read the introduction to the present volume you will have met the phrase already, because it is the frame on which everything here hangs. The axemaker's gift is the bargain struck by that first person who chipped a

cutting edge and handed it round: the tool gives you powers you never dreamed of, and the price is that the world changes in ways nobody can foresee, and there are no refunds. We told the story of the axe, and the plough, and the printing press, and the bomb, and we argued that the very same cast of mind that could make a fine cutting tool — the mind that cuts, that separates, that takes the seamless world and slices it into manageable, nameable, orderable pieces — is the mind that made everything else, for good and ill. The axemaker cuts stone. Then he cuts wood, and cloth, and eventually the atom, and the working day, and the enemy. The axe is not just a thing. It is a *way of thinking*, and the way of thinking is this: to control something, first cut it into pieces small enough to handle, and name them.

Language is the axe turned on the world itself. This is the master move, the one from which all the others descend. A wolf experiences the forest; a human being names the forest, and having named it can now name the trees, and the parts of the trees, and the parts of the parts, and — crucially — can hold those names in the mind and shuffle them about *in the absence of the things themselves*. This is the trick that changes everything. The word "fire" is not hot. You can say it in the dark and the cold. You can say "fire" to a child who has never seen one and, if you are patient and the child is clever, plant in that small head a working model of a thing not present — a stand-in, a token, a portable ghost of the real. And once you can do that, once a sound can stand in for a thing, you are no longer confined to the here and the now and the visible. You can talk about yesterday. You can talk about the far valley. You can talk about the dead, and the not-yet-born, and the herd that crosses in autumn, and the god behind the mountain, and things that have never existed and never will. The animals live in the present tense. We are the creatures who escaped it, and we escaped it through words.

There is a very old story about this, the oldest we have, and I put my epigraph at the head of this chapter for a reason. In the garden, before the

fall, before the woman, before anything much happens at all, God does a strange and rather touching thing: he makes the animals, and then he brings them to the man, one by one, to see what he will call them. Not to see whether the man will get the names *right* — there are no right names waiting to be guessed — but simply to see what he will call them, and whatsoever Adam called each creature, *that was its name*. It is the first job God gives to humanity, before tending the garden, before anything: the naming of the parts. And it is a job that confers a peculiar authority, because in the old way of thinking to know the true name of a thing was to have power over it — which is why the demons in the exorcisms must be made to speak their names, and why Rumpelstiltskin comes apart the moment the queen guesses his. Whoever wrote down that scene in the garden, three thousand years ago and more, understood something we have only lately built a machine to prove: that the naming of the world and the mastering of the world are, at some deep level, the same act. To name is to seize. Adam does not merely observe the animals. He takes possession of them, with his mouth.



But I have got ahead of myself, because for a very long time — far longer than we comfortably admit — all of this took place entirely in the air. The word was only ever spoken, heard, and gone. There was no writing, would be no writing for the overwhelming majority of the human career, and everything a people knew — every skill, every law, every story, every genealogy and prayer and the location of every reliable waterhole for a hundred miles — had to be carried in living heads and passed from mouth to ear or lost forever. This is the condition we call orality, and it is so unlike our own bookish existence that it takes a real effort of imagination to inhabit it, but the effort is worth making, because the machine at the end of this book is, in one uncanny respect, a throwback to it.

Consider the problem. If everything must be remembered, and there is no way to look anything up, then memory is not a party trick; it is the load-bearing wall of the entire culture. Lose the old woman who knows the songs and you have burned a library, except that nobody can smell the smoke. So the oral cultures became, of necessity, prodigious rememberers, and — more interesting still — they shaped their most precious knowledge into forms *built to be remembered*. This is why the foundational texts of so many peoples come to us as poetry. Rhythm, rhyme, metre, repetition — these are not decoration. They are error-correction. They are the format in which a long message can be held stable in a human head across a lifetime and handed on without drifting into nonsense, because the metre itself tells you when a word has gone missing, the way a gap in a paved floor tells you a tile has fallen out.

The great demonstration of this — and it is one of the loveliest pieces of detective work in the humanities — came from a young American scholar named Milman Parry in the 1920s and 30s. For two and a half thousand years people had argued about Homer: how could one blind poet, in an age before writing had any grip in Greece, have composed and held in his head the whole vast sprawl of the *Iliad* and the *Odyssey*? Parry noticed something everyone had seen and nobody had understood — that Homer is full of stock phrases, the same epithets endlessly recycled: the wine-dark sea, the rosy-fingered dawn, swift-footed Achilles, the grey-eyed goddess Athene. Generations of schoolmasters had treated these as a sort of quaint tic, or padding. Parry realized they were *machinery*. Each formula was metrically shaped to slot into a particular position in the line, so that the singer, composing at speed out loud in front of an impatient audience, was not remembering the poem word for word at all — he was rebuilding it, live, each time, out of a vast kit of ready-made, metre-fitting parts, the way a jazz musician improvises a solo out of learned phrases. And then Parry did the thing that makes him a hero of mine: he went to check. He travelled to what was then Yugoslavia, into the mountains, where illiterate

peasant bards — the *guslari* — were still, in the twentieth century, singing epic poems of tremendous length to the scrape of a one-stringed fiddle, composing them anew at every performance by exactly the method he had reconstructed from Homer. He recorded them. He proved it. Homer was not a man who had memorized a book. Homer was a *technique*, a way of generating fluent, structured, apparently spontaneous speech from a stock of learned patterns.

I ask you to file that description somewhere safe, because we will want it again before the end. A system that produces fluent, plausible, well-formed language on demand, not by looking up a fixed text but by recombining an enormous store of learned patterns into whatever the moment requires. If that begins to sound familiar, you are paying attention, and we are only in the first chapter.



Now, the trouble with the spoken word — its glory and its curse — is that it vanishes. It exists only in the instant of its saying, and then it is gone into the air, and all you have left is what your fallible, forgetful, wishful head chose to keep. For a hundred thousand years, give or take, this was simply the human condition, and it had consequences we can still feel in our bones. An oral culture cannot easily accumulate. It can conserve, magnificently — the old songs, the old skills, held with a fidelity that shames us — but it struggles to *build*, because each generation must spend so much of its capacity just holding on to what it already has. There is no filing cabinet. There is no reliable outside store. The knowledge of the tribe can never much exceed what its living members can between them carry, and when they die, unless they have passed it on entire, it dies with them.

The escape from that trap is the second great abstraction in this story, and like the first it came, gloriously, from nobody trying to do anything of the sort. It came from accountants.

I love this, and I want to slow down for it, because it is so perfectly, characteristically the way real change happens — sideways, by accident, in the pursuit of something entirely mundane. For most of the last century the textbooks told a heroic story about the birth of writing: that some Sumerian genius, gazing at the stars or the gods, conceived the sublime idea of capturing speech in visible marks. It is a lovely story and it is almost certainly wrong. The truth, painstakingly assembled by an archaeologist named Denise Schmandt-Besserat, who spent years doing the unglamorous work of cataloguing little lumps of baked clay in museum drawers across the Near East, is far more interesting and far more human. It turns out that for thousands of years before writing, all across the ancient Middle East, people kept accounts using small clay tokens — little cones, spheres, discs, cylinders, each shape standing for a quantity of something: this cone is a small measure of grain, this sphere a large one, this disc a sheep. If I owed you five sheep and three jars of oil, we would settle the debt with the right tokens. Simple. And utterly ordinary — the boring bookkeeping of an economy learning to run granaries and temples and taxes.

Then, around five and a half thousand years ago, somebody had the bright idea of sealing the tokens for a given transaction inside a hollow clay ball — a *bulla* — so the count could not be tampered with in transit: your obligation, locked in a clay envelope, witnessed and secure. Excellent. Except now you had a new and idiotic problem, which is that you had sealed the tokens inside an opaque ball and could no longer see what was in it without smashing it open. So — and here is the hinge of world history, disguised as an office-supplies annoyance — some clerk began pressing the tokens into the soft surface of the ball *before* sealing them inside, so the outside would show what the inside contained. Three sheep-tokens pressed into the clay: three sheep-marks on the shell.

And then someone, some blessed, lazy, immortal someone whose name we will never know, looked at the marks on the outside and looked at the

tokens on the inside and had the thought that made the modern world: *why do we need the tokens at all?*

If the marks say three sheep, keep the marks and throw away the sheep-tokens. Keep the *record* and discard the *thing*. Draw the sign for sheep on a flat tablet and be done with it. And with that shrug — that supremely practical, supremely unvisionary shrug — writing was born, not as a gift from heaven but as a shortcut in a warehouse. The first written documents in human history are not poems or prayers or the deeds of kings. They are receipts. Inventories. So-and-so owes the temple this many measures of barley. The word made visible began its career as a spreadsheet.



But look what they had done, these Sumerian clerks, without the least idea they were doing it. They had performed the second great abstraction, and it is worth naming as precisely as we can, because it is the very move the talking machine will one day perform on all of us. The first abstraction, remember, was the spoken word: a sound stands in for a thing. The fire-word for the fire. The second abstraction is a mark that stands in for the word that stands in for the thing — a symbol of a symbol, reality at two removes, the ghost of a ghost. And once you have that, once meaning has been successfully peeled off the fragile, mortal, forgetful human head and fixed in durable clay, the trap of orality springs open. Knowledge can now sit outside the mind. It can outlive the knower. It can accumulate, tablet on tablet, generation on generation, into something no single head could ever hold. It can be sent down the river to a man you will never meet and understood exactly as you meant it. The word has been made to keep.

The Sumerians wrote by pressing a cut reed into wet clay, which leaves a wedge-shaped stroke — *cuneus* is Latin for wedge, hence cuneiform — and at first the signs were little pictures: a stylized ox-head meant ox, a stalk meant grain. But pictures can only take you so far, because most of the things worth saying are not picturable. How do you draw *tomorrow*, or

debt, or *swear*, or the name of a man called, say, Enlil? And so the Sumerians stumbled onto the third and cleverest trick of all, the one that finally welded writing to speech: they let the signs stand not for things but for *sounds*. It is the principle every child discovers in the puzzle-page where a picture of an eye plus a picture of a saw means "I saw" — the rebus. The sign for a thing whose name sounded like part of your word could be borrowed to spell that word, regardless of meaning. And at that instant writing stopped being a system of pictures of the world and became a system for *capturing the noises a mouth makes* — which meant, at last, that anything that could be said could be written, and the two channels, the spoken and the written, locked together into the double helix that has carried human knowledge ever since.

It was still, mind you, fiendishly hard. Cuneiform had hundreds of signs and took the better part of a decade to master, which meant that the men who could read and write — the scribes — held a monopoly on the frozen word, and a monopoly, as we shall see again and again in this book, is a thing people fight to keep. But the difficulty is worth dwelling on, because the next great step is the step that broke it.



Somewhere in the Bronze Age, among Semitic-speaking people on the fringes of the Egyptian empire — the earliest traces turn up scratched in the turquoise mines of the Sinai, so we call it Proto-Sinaitic, and they were very possibly labourers, not scholars — somebody had the most consequential idea in the history of literacy. Instead of a sign for every word, or even a sign for every syllable, why not a handful of signs, one for each basic *sound* in the language? A sound so small it cannot be divided — what we would call a consonant. Two dozen marks or so, and with them you could spell anything at all. They built the signs, charmingly, out of little pictures whose names began with the sound they wanted: a picture of an ox — '*aleph* — for the sound that starts it; a picture of a house — *bet* — for the

b; and so, in time, down the whole line, until you had a fixed running order of little sound-pictures that a child could learn in an afternoon. We still recite them. We just say the first two of them run together, and call the whole thing after them, and the word is *alphabet*. Every A you have ever written is a picture of an ox, upside down; turn a capital A over and the two legs become the horns. You are drawing cattle and don't know it.

The Phoenicians, those tireless Mediterranean traders, carried this consonant-alphabet the length of the sea, and the Greeks, borrowing it, made one further adjustment that the schoolbooks tend to skip over but that mattered enormously: they took a few of the consonant-signs they didn't need for Greek and repurposed them to mark the *vowels*, the open sounds a Semitic writer had left the reader to supply. It sounds like a footnote. It was a revolution, because a fully phonetic alphabet — every sound, consonant and vowel alike, set down explicitly — is a machine so simple and so complete that a person can learn to read it without a teacher standing over them to fill in the gaps. Writing had spent three thousand years as the closely guarded mystery of a scribal priesthood. The alphabet was the first crack in that wall — the first time the frozen word came within reach of the ordinary person — and it is surely no coincidence at all that the Greeks who inherited it went on, within a few centuries, to invent so much of what we mean by argument, and philosophy, and history, and democracy: for the first time, the tools of the word were not locked in a temple. But that is a thread that runs through the whole of what follows, and I will not pull it here.



Let me stand back, because we have come a long way in a few thousand words — from a stone too big to hold to a set of two dozen marks that can hold anything — and I want to name the single deep tendency that runs through the whole of it, because it is the tendency that, compounded across a hundred centuries, finally wakes up.

At every step, the story is the same. We take something continuous, seamless, whole — the flow of an action, the blur of the world, the stream of speech — and we *cut it into parts*, and we *name the parts*, and we *set the named parts in order*. The knapper cuts the making of the axe into a sequence of blows. Language cuts the world into a sequence of words. Writing cuts speech into a sequence of visible signs. The alphabet cuts those signs down to the irreducible atoms of sound. Down, and down, and down — always the same axemaker's motion, the cut that separates so that the pieces can be handled and named and rearranged. And this, this habit of taking things apart to see how they work and describing them piece by piece, is the deepest groove in the whole of the Western mind. It is the thing I once called, in another book, a tool for chipping pieces off the universe. The philosophers have a grand name for it — reductionism — and out of it, in the seventeenth century, science itself would be born, the systematic business of dismantling reality into ever smaller, ever more precisely named components: matter into elements, elements into atoms, the atom into its own private zoo. But you do not need a laboratory to see it. You need only a mouth, and the sequential, cutting, naming habit that the mouth learned, perhaps, from the hand that chipped the stone.

The purest monument to this cast of mind — and one of my favourite characters in the whole gallery — is a Swedish parson's son called Carl Linnaeus, who in the eighteenth century looked at the whole of the living world, which arrives, you will recall, as a great blooming buzzing confusion, and decided to *file* it. Every plant, every animal, every living thing, sorted into a great nested hierarchy of kingdoms and classes and orders, and each one given a two-word Latin name — a genus and a species — so that a rose was *Rosa*, a particular rose *Rosa canina*, and a naturalist in Uppsala and a naturalist in Philadelphia who shared not one word of any living language could nonetheless know they were discussing precisely the same creature. It was Adam's task in the garden, resumed after several thousand years and carried through with Nordic thoroughness. Linnaeus

named some twelve thousand species, and when he came to the naming of the creature doing the naming, he did not flinch: he reached for the flattering word, and called us *Homo sapiens*, man the wise, which is either a magnificent piece of self-confidence or a very early example of the machine marking its own homework. There is a line attributed to him — probably by an admirer rather than the man himself, so take it with the usual pinch of salt — that ran *Deus creavit, Linnaeus disposuit*: God created, Linnaeus arranged. And whether he said it or not, it is exactly right, and it is exactly the human story: we did not make the world, but by God we named it, and in naming it we took, as Adam did, a kind of ownership. Linnaeus did not add a single creature to creation. He added *words*, arranged in order, and the words let us hold the whole living world in a way no one before ever had. A machine for arranging names had reorganized life itself.



And now I can state the thesis of this book plainly, because everything so far has been building the ramp up to it, and if I have done my job you may find you have half-guessed it already.

Here it is. For a hundred thousand years and more we have been perfecting a single technology — the arranging of words — and in the course of perfecting it we have quietly done something we have never quite admitted to ourselves. We have moved almost everything that matters to us *out of the world and into words*. We began by describing reality; we ended by living inside the description. And the reason this matters — the reason it is the whole hinge of the age we have just entered — is deceptively simple: if the important things are made of words, then a machine that can do nothing whatever except arrange words can, in principle, do the important things.

Watch me make good on that, because it sounds like a cheap trick and it is not.

Take the law. What, physically, is a law? Where is it? You cannot stub your toe on the law of contract. A statute is words — words arranged in a particular authoritative order, and nothing else. A contract is words. A verdict is words, a constitution is words, a marriage is a few dozen words said in the right order in front of the right person, a nation's sovereignty is words on a treaty, a man's guilt or innocence is a word pronounced by twelve other people. The entire majestic apparatus of the law — which can take your money, your children, your freedom, your life — has no material existence at all. It is a structure built entirely out of language, sustained entirely by our collective agreement to treat certain arrangements of words as binding. Adam named the animals; the lawyer names the crime, and the naming makes it so.

Take money — and here the point gets sharper, because we imagine money is a thing, gold and coin and heft, and it has not been for a very long time. Look at an old English banknote and it tells you, in so many words, exactly what it is: *I promise to pay the bearer on demand*. A promise. Words. The overwhelming majority of the money in the world today is not notes and not coin and certainly not gold; it is numbers in the ledgers of banks, entries in a database, a shared and updated agreement about who owes what to whom — which is to say it is, almost in its entirety, a very large and very serious arrangement of symbols that we have all consented to take seriously. Money is a story we tell so consistently that it comes true. It is the most powerful fiction ever composed, and it is made of exactly the same stuff as the *Odyssey*: words in an agreed order.

Take science — the very thing that the cutting, naming, reductionist habit was invented to do. What does a scientist actually produce? Not, in the end, a thing. A *paper*. A description. Words about measurements, laid out in a precise and testable order: under these conditions, this was observed; therefore we propose this. The knowledge is not in the phenomenon; the phenomenon was always there. The knowledge is in the *description*, and the description is language — sometimes the specialized

language of mathematics, which is only language with its collar buttoned, but language all the same. Newton did not add gravity to the universe. He added an equation — a sentence in the dialect of number — that let us hold the falling of every apple and the turning of every planet in a single graspable phrase. The world does the falling. We do the wording.

And take — because we must, since it is the whole point of what is coming — the machine's own native realm, code. What is a computer program? People imagine something electric and alien, and underneath, yes, there are electrons doing incomprehensible things at incomprehensible speed. But a program, the thing a human being actually writes, is *text*. It is words — words with a punishingly strict grammar, admittedly, because the reader is a machine with no gift for guessing what you meant, but words in an order, sentences of instruction, and nothing more. Code is not merely *like* language. Code *is* language: a way of arranging symbols such that the arrangement, read out, makes something happen. The software that runs the bank that holds the money that pays the lawyer that argues the law — all of it, top to bottom, words upon words upon words, all the way down.

Do you see it now? Law, money, science, software, government, scripture, scholarship, the whole towering scaffolding of civilization that we imagine to be so solid — nearly the whole of it turns out, on inspection, to be *made of language*, built out of the one abstraction we have been compounding since somebody sat by a fire and taught a child to say *strike here, not there*. We did not notice, because it happened over ten thousand years and we were busy. But we have spent the entire span of history migrating the important business of being human out of the physical world and into a world made of words — a world you cannot touch, that exists only in the saying and the writing and the shared agreement to go on treating the symbols as real.

And that migration is why the machine at the end of this book is not a gadget. It is the punchline of the whole hundred-thousand-year joke.

Because we have finally built a device that does one thing and one thing only: it arranges words. It has no hands, no eyes in any human sense, no body, no world; it has never seen a sheep or felt a fire or owed anybody five measures of barley. It does nothing but predict, with uncanny fluency, what word should come next — the *guslar's* trick, Homer's trick, done with electricity and at a scale that would have stopped Milman Parry's heart. And because we had already, without meaning to, poured nearly everything we care about into the medium of words, a machine that is fluent in words turns out to be, willy-nilly, fluent in *nearly everything*. It can draft the contract because the contract is words. It can write the code because the code is words. It can explain the science because the science is words about measurements. It can argue the case, compose the prayer, tell the story, name the animal. It inherited the estate simply by learning to speak, because we had already put the estate into speech.



I want to be careful here, at the hinge, not to claim too much, because overclaiming is how technologists make fools of themselves and I have watched a great many do it. A machine that arranges words is not thereby wise, and it is certainly not reliable — I have personal reasons to know this, having been informed by one of these very machines, in impeccable prose and with total confidence, that I died peacefully in 2019, which was news to me and, I should think, to my publisher. The machine's fluency and the machine's truthfulness are two entirely different things, and the gap between them is going to occupy us badly in a later chapter, so I will only wave at it now. The point is not that the word-machine is right. The point is that it is *fluent*, and that fluency, in a civilization that has spent ten millennia encoding itself in words, is a skeleton key.

Nor — and this matters — is any of this quite as new as the newspapers shrieking about it seem to think. That is really the burden of my whole song. We have met the anxiety before, every single time, at every one of the

steps I have just walked you through. When writing arrived, the complaint (recorded, exquisitely, by Plato, in a passage we will come back to) was that it would rot the memory and give men the mere show of wisdom without the substance. When the alphabet democratized the frozen word, the guardians of the old scribal mystery were no more pleased than any guild ever is to see its monopoly broken. Every extension of the word-tools has provoked exactly this mixture of wonder and dread, and every time the same two things have turned out to be true: that the critics were right about what would be lost, and that the change came anyway, and that on the far side of it we found powers we could not have imagined and could not, afterwards, imagine living without. The talking machine is the newest turn of a very old screw. It is the axemaker's gift arriving, as it always does, with the powers gift-wrapped and the price in the small print and no receipt.

The one genuinely new thing — and I say this so that you do not think me merely a debunker, waving away the marvel with a historian's tired hand — is the *closing of the loop*, and it is worth pausing on, because it is strange enough to keep you up at night if you let it. Every previous word-tool was a tool for *handling* words: the mouth made them, the clay kept them, the alphabet simplified them, the press — but I am trespassing on a later chapter — multiplied them. In every case the words were about the world, and we, the word-users, stood outside the system, working it. What is different now, for the first time in the whole long story, is that we have built a tool that does not merely store or carry or copy our words but *generates* them — that sits, as it were, in the chair the human being has occupied alone since the fire and the flint, the chair of the one who decides what word comes next. And we built it, extraordinary to relate, by feeding it our words. All of them, or as near to all as makes no difference: the receipts and the epics, the statutes and the love letters, the science and the scripture and the lies, the whole ten-thousand-year sediment of everything anybody ever bothered to write down, wound onto a loom and woven into

a single vast fabric from which the machine now spins its fluent, plausible, sometimes entirely fictional continuations of the human sentence. It is made of us. Every word it says, it learned from a word one of us said first. When it told me I was dead, it was only repeating, in its patient average way, a rumour it had assembled out of us — for somewhere in that great woven web of our words, no doubt, some careless human had written that James Burke was no longer with us, and the machine, faithful as a mirror and about as wise, wove it back.



So here is where we stand, at the end of the beginning, with the rest of the book stretching out ahead like a trail of clues.

A million and a half years ago, more or less, one of our ancestors made a stone too fine for its job, and in the making of it — in the precise, sequential, teachable business of striking flakes in the one right order — there lay, some say, the seed of language: the sequential, part-by-part, name-the-pieces habit of mind that would eventually take the whole world apart and lay it out in a line, word by word by word. That habit gave us speech, and speech let us carry the world about inside our heads and speak of the absent, the past, the future, the never-was. Then, chasing nothing grander than an accountant's tidy mind, we froze the word in clay, and it began to keep and to pile up and to outlive us; and we cut the frozen word down to two dozen sounds a child could learn, and it slipped the grip of the priesthood and spread; and all the while, without ever deciding to, we were pouring the whole serious business of human life — the law, the money, the science, the promises, the very machinery by which we hold society together — out of the touchable world and into the untouchable world of words. Until at last there was so much of us stored in words, and our tools for handling words had grown so powerful, that we did the thing that closes the circle. We took everything we had ever said, and we taught it to speak.

Everything that follows in this book is the story of how the word-tools compounded, step by unwitting step, until they woke up. It is a story with no villains and no masterplan, only a very long chain of people solving the small practical problem in front of them and having not the faintest notion what they were building — a silk-weaver, a census clerk, an artilleryman, a man watching a dead frog twitch, a telephone engineer counting the redundancy in the English language, a pair of programmers making video-game explosions look wetter. Not one of them thought they were making a mind. Neither, I rather suspect, do the people making it now. But that is the way of the axemaker's gift, and always has been, and the deepest lesson a historian has to offer is also the most unsettling and the most consoling at once: that the future is only the past, waiting to happen. We have stood exactly here before. Not with this tool — but with this moment, the moment when a new way of handling words breaks loose and rearranges the world. It went a certain way each of the previous times, and I can show you how, and I intend to, in the chapters that follow.

But to do that we must go back to the beginning of the *machinery*, which is not, as you might expect, a laboratory or a war room, but a workshop in the French silk trade, full of angry weavers and beautiful cloth, where a man was about to solve a problem about flowered patterns and, in solving it, hand the future a way to remember instructions on a string of punched cards.

Which is where, as we shall see, the thread is first drawn onto the loom.

CHAPTER TWO

Weaving the Thread

*We may say most aptly, that the Analytical Engine weaves
algebraical patterns just as the Jacquard-loom weaves flowers and
leaves.*

— ADA, COUNTESS OF LOVELACE, NOTE A TO HER
TRANSLATION OF MENABREA'S *SKETCH OF THE
ANALYTICAL ENGINE*, 1843

On Saturday evenings in the 1840s, at number one Dorset Street, Marylebone, you could meet everybody. Charles Babbage's soirées were a fixture of the London season, and over the years the guest lists read like the index of the age — Darwin, Dickens, Brunel, the Duke of Wellington, foreign ambassadors, bishops, actresses, and a steady supply of duchesses. They came for two attractions. One was the Silver Lady, a beautiful automaton dancer Babbage had first seen as a boy in a mechanical museum, had bought back at auction decades later and restored with his

own hands, and who pirouetted, fluttered a clockwork bird, and was mobbed all evening. The other was a gleaming brass-and-steel something in the next room, which could do sums. Babbage noted sourly that his English guests crowded round the dancer while the occasional foreign visitor stood alone, thinking hard, in front of the machine. He minded about that more than he ever admitted.

But the object I want you to look at hung quietly on the wall: a framed portrait of a seated gentleman in a comfortable coat, holding a pair of dividers, his workshop arranged around him. Babbage liked to ask new arrivals what they made of it, and the guests would glance at it and pronounce it a very fine engraving. It was nothing of the kind. Put a magnifying glass to it — Babbage would happily hand you one — and the picture dissolved into thread. It was silk, woven in Lyon, every shadow of the man's face formed by the crossing of warp and weft, and the loom that wove it had been instructed, thread by thread and row by row, by roughly twenty-four thousand punched cards strung together in sequence. The comfortable gentleman in the picture was Joseph-Marie Jacquard, the man whose cards they were. Which means the thing on Babbage's wall was, if you'll allow the anachronism, a portrait of a programmer, generated by his own program — and the host who kept pointing it out to baffled guests was at that very moment designing a machine that would take the same punched cards and use them to weave, as his interpreter would put it in the sentence at the head of this chapter, patterns of algebra instead of patterns of flowers.

Nobody at the party got the point. Parties are like that. But you and I have the advantage of knowing how the story comes out, so let me tell it properly. It begins, as stories about information so often do, with somebody stuck in a truly wretched job.



Silk was in Lyon because the kings of France were tired of watching good money leave for Genoa and Florence every time the court needed something splendid to wear; in the 1530s the city was granted privileges to make the stuff itself, and over two centuries the trade grew into *la Grande Fabrique*, the great fabric, which at its height set tens of thousands of looms clattering in the tall houses of the city. And what Lyon sold, above everything, was the *figured* silk — brocade, damask, lampas — cloth with pictures in it: flowers, birds, pomegranates, whole gardens rippling across a court dress. Plain silk is merely expensive. Figured silk was power made visible, and Europe's aristocracy could not get enough of it.

Now, how do you get a picture into cloth? On a plain loom the weaver raises half the warp threads, throws the shuttle through the gap, lowers them, raises the other half, throws it back. Over, under, over, under. But a picture requires that on every single pass of the shuttle, some *particular selection* of warp threads — this one, not that one, these forty, not those sixty — be lifted, and a different selection on the next pass, and the next, for thousands of rows. The instrument that did this was the drawloom, an idea as old as China, and it needed two human beings. The weaver sat at the front and wove. And up in the wooden framework above the warp perched a child.

The child was called the drawboy, and I would like you to spend a moment in his shoes, because he is the reason this chapter — and in a sense this book — exists. The pattern of the fabric had been translated, by a specialist who "read" the design off squared paper, into an arrangement of hundreds of looped cords. The drawboy's job was to pull, on the weaver's shout, exactly the cords that lifted exactly the warp threads needed for that one row of the picture — haul, hold against the tension of the loom while the shuttle went through, release, find the next set, pull again. All day. In a good day the pair of them might produce a couple of inches of brocade, or rather less if the drawboy's attention wandered, one wrong cord being enough to plant a flaw in the flowers that would catch the light forever. It

was tedious, exhausting, deforming to growing bodies, and it paid the way jobs pay when the labour is supplied by children. Adam Smith would have recognized the arrangement instantly; so, I'm afraid, would the manager of a modern data-labelling office.

But notice — because the whole story turns on it — what those cords actually were. The design existed in three forms. It was a painting on squared paper in the studio. It was a piece of shimmering cloth on the loom. And in between, it was *the arrangement of the cords*: a physical, storable, consultable encoding of the pattern, read out one line at a time by a small boy acting as the playback mechanism. Lyon had built, without ever thinking of it this way, a memory made of string — and had staffed the retrieval system with children. The pattern was information. The information was the valuable part; designs were commissioned, hoarded, copied and stolen, because in the fashion trade — and Lyon lived and died by fashion, by the season's novelty, by the fact that this year's pomegranates were worthless the moment next year's tulips appeared — the cloth was merely the printout. Anyone who found a way to store the pattern more cheaply, and play it back without a drawboy, was going to be a very popular man.

Or, as it turned out, a very unpopular one.



The first person to see the answer was a weaver's mechanic named Basile Bouchon, and he saw it because of what his father did for a living. Bouchon senior built organs. Automatic music was already a flourishing trade — carillons in Flemish belfries played tunes from great pinned drums, and the barrel organ stored its melody as a pattern of metal pins on a rotating cylinder, each pin a note, each blank a silence. Music, in other words, had already been digitized — a tune reduced to a sequence of discrete yes-or-no events, stored in durable form, playable by anyone who could turn a crank. The younger Bouchon, surrounded by looms all week and pinned barrels

on Sundays, did the obvious thing that nobody had done: in 1725 he rigged a loom so that a roll of paper, punched with rows of holes, pressed against a row of needles. Where there was a hole, the needle passed through and its thread was selected; where the paper was blank, the needle was pushed back and its thread stayed put. Hole, lift. No hole, don't. The pattern had moved from cords to paper, and the drawboy's judgement had been replaced by perforation.

It half worked, which in the history of technology is the most dangerous amount of working. The paper tore, the alignment drifted, and the loom still needed an assistant. Three years later a colleague named Jean-Baptiste Falcon made the improvement that matters more than it looks: he replaced the fragile roll with a chain of stiff punched *cards*, laced together with cord. Cards could survive the workshop. Cards could be corrected — punch a new one, lace it in where the error was, throw the bad one away. Cards could be duplicated, filed, sold, and strung into sequences as long as you pleased: a complete anticipation of the first principle of software maintenance, namely that you will get it wrong the first time and would prefer not to start over. A few dozen of Falcon's looms found buyers. Most of Lyon shrugged. The machines were finicky, the drawboys were cheap, and the merchants who ran the Fabrique had no pressing reason to care how the weavers' children spent their childhoods.

Enter, in the 1740s, one of the strangest figures in this whole strange story: Jacques de Vaucanson, celebrity engineer, showman, and creator of the most famous fraud in the history of automata. Vaucanson had made his name in Paris in 1738 with a mechanical flute-player that genuinely played the flute — lips, tongue, fingers, breath, the whole embouchure done in machinery — and he followed it a year later with the sensation of the age: a gilded copper duck, life-sized, which took grain from your hand, swallowed it, and then — the part that emptied theatres of their composure — visibly digested it, producing from the appropriate end a convincingly authentic result. Voltaire remarked, with a straight face that

has kept scholars guessing ever since, that without Vaucanson's duck there would be nothing to remind you of the glory of France. And the digestion? A century later the conjuror Robert-Houdin — from whom Houdini would borrow a name — was hired to repair the ageing bird and found inside a hidden compartment pre-loaded with green-dyed breadcrumb paste. The duck had never digested anything. The most celebrated demonstration of mechanical life in the eighteenth century was, at the business end, a man palming the evidence. I mention this partly because it is funny, and partly because it is the earliest and best illustration of a law we will need again and again in this book: whenever a machine astonishes you with how lifelike it is, check the other end for breadcrumbs.

But Vaucanson was no mere showman, and the French state knew it. In 1741 he was appointed inspector of the silk manufactures, took one look at the drawloom, and did what a first-rate mechanic does: he abolished the boy. His loom of 1745 put the pattern-reader on top of the frame, drove it automatically from the loom's own motion, and selected the threads from a perforated cylinder. It was the most advanced textile machine on Earth, and Lyon hated it. The weavers saw, quite correctly, what it meant for the trade as they knew it, and they rioted; Vaucanson was pelted with stones in the street, and — the story goes, and I pass it on at the price I paid for it — took his revenge by building a loom worked by a donkey, so that all Lyon could see flowered silk being woven by an ass. The city declined to be amused or converted. The loom was crated off to Paris, where it eventually came to rest in the Conservatoire des Arts et Métiers, the great national attic of machines, and gathered dust for half a century. Remember that dust. A museum, it turns out, is not where ideas go to die. It is where they go to wait.

★

The man they were waiting for was born in Lyon in 1752, into the trade and at the bottom of it. Joseph-Marie Jacquard was a master-weaver's son,

and the accounts agree that as a boy he was put to the work the sons of weavers were put to — including, they say, his time in the cords as a drawboy, which he loathed with the specific, patient loathing that produces inventors. He tried bookbinding, type-founding, cutlery; he inherited looms and lost them; and then history ran over him, as it ran over everyone in Lyon in the 1790s. The city rose against the Convention in 1793 and was besieged, taken, and punished with a ferocity remarkable even by the standards of that decade. Jacquard fought — first, it seems, in the defence of the city, then, prudently rebadged, in the armies of the Republic, side by side with his teenage son, who was shot dead beside him. He came home to nothing very much and went back to the looms. If you had met him in 1799 — a middle-aged, twice-ruined ex-soldier restoring old looms in a half-destroyed city — you would not have picked him as the man about to hand the future its operating principle.

What pulled him out of obscurity was, of all things, an English prize. A society in London had offered a premium for a machine to make fishing nets, the challenge crossed the Channel, and Jacquard, tinkering, produced one. It got him noticed, and in 1803 it got him summoned to Paris, to the Conservatoire — where, set to work among the national collection of mechanical marvels, he came face to face with Vaucanson's dusty loom. And here the threads of half a century braided themselves together in one workshop: Bouchon's punched holes, Falcon's laced cards, Vaucanson's automatic mechanism, examined and combined by a man who had personally been the component all three were designed to replace. By 1804 Jacquard had his patent, and the machine the world would call by his name — though you will have noticed by now that, like most eponymous inventions, it contained rather a lot of other people — sat on top of an ordinary loom and did the whole job: an endless chain of punched cards advancing automatically, one card per throw of the shuttle, each card selecting, through spring-loaded hooks and needles, exactly the threads to rise for that row. One weaver, working alone, could now produce figured

silk faster than a weaver-and-drawboy team, to any design whatever, of any length whatever. Change the cards and the same loom wove a different picture. Punch a fresh chain and it wove a picture nobody had ever seen.

Stop on that sentence, because the whole twentieth century is folded up inside it. *Change the cards and the same machine does something else.* Every machine humanity had built until that moment had its purpose in its bones — a mill grinds, a pump pumps, and if you want a different job done you build a different machine. The Jacquard loom split the machine in two: an unchanging body of iron and wood, and a changeable soul of punched cardboard. The body was capital; the soul was information. The soul could be stored on a shelf, copied by apprentices, sold by the pattern-house, pirated by competitors, posted to Vienna. Lyon's design studios were, though the word was a century and a half away, software houses — rooms full of specialists translating human desires into machine-readable code, one hole at a time, for an industry whose product was obsolete every season. And that is why I insist, against the whole dignity of my profession, that software began as fashion. Not in the laboratory, not in the counting-house, not in the ministry of war: in the rag trade. The first industry to store its soul in punched holes was the industry of looking marvellous. We taught cardboard to remember, so that a duchess should never wear last year's flowers.

Lyon's weavers greeted the future with their usual perceptiveness: they tried to kill it. The machine was denounced, and Jacquard was mobbed; the *conseil des prud'hommes*, the trade's own tribunal, is said to have had one of his looms broken up in public, the iron sold as iron and the wood as wood, and there is a persistent story that its inventor was nearly thrown into the Rhône by a crowd of weavers. I owe you a confession about that last item: it comes to us polished by Samuel Smiles, the Victorian apostle of self-help, who never met an inventor's tribulation he couldn't improve, so treat the drowning as decoration and the fury as fact. The fury was entirely real, and — as we shall see when this book reaches Nottingham, where men would

soon hang for the English version of the same rage — it was not stupid; but that argument belongs to Chapter Eleven, and I shall keep it there. What matters here is that the machine won, and won quickly, because this time the state was buying. Napoleon — at war with Britain, and acutely aware of what Britain's mills were funding — came to Lyon in 1805, inspected, and understood. The patent was bought for the public: the design was handed to the city of Lyon, and Jacquard exchanged his monopoly for a pension of three thousand francs and a royalty on every loom. Within a few years the machines were spreading through the city that had tried to break them — some accounts say eleven thousand in France by 1812 — and the silk they wove was cheaper, faster, and as gorgeous as anything the ancien régime had ever draped over itself. Jacquard lived to be old and honoured; the city that had allegedly ducked him in the Rhône eventually put up a statue; and the drawboys went — some into the mills, some to school, and the job itself into the museum of abolished miseries, where I hope it is displayed next to a placard reading *nobody missed it*.

And in 1839, five years after the old man's death, the Fabrique produced its own monument, and a very Lyonnais monument it was: not marble but silk. In the workshop of Didier Petit et Cie, a master weaver named Michel-Marie Carquillat took a portrait of Jacquard painted by Claude Bonnefond and translated it — or rather, had it translated, hole by hole, into that chain of some twenty-four thousand cards — and wove it, in black and white silk, at a fineness that fooled the eye of everyone who ever glanced at it casually. Copies went out, whenever a dignitary needed impressing, to the great and the curious across Europe. One of the curious was an Englishman who understood, better than any man alive, exactly what he was looking at. He hung it in his drawing room in Dorset Street and began asking his party guests what they made of it.



Charles Babbage was rich, brilliant, generous, quarrelsome, and incapable of finishing things — a combination the reader may recognize from other walks of life. He held Newton's old chair of mathematics at Cambridge for eleven years without, as far as the records show, delivering a single lecture, and he compensated by doing everything else: he helped found the Astronomical Society and the Statistical Society, invented the cowcatcher, showed that the cost of a letter lay in the handling and not the distance (an argument that fed the penny post), broke ciphers, investigated volcanoes, and campaigned against organ-grinders with a venom he never quite managed to direct at his real enemies. But the work of his life began, by his own account, one evening around 1821, when he and his friend John Herschel sat checking astronomical tables that had been computed twice, independently, by hired human calculators. The two sets of results disagreed with each other, and with the truth, in error after error, and Babbage, head in hands, said: "I wish to God these calculations had been executed by steam."

It is worth pausing to feel the force of that wish, because tables are the least glamorous objects in this book and they ran the world. Every ship that found its longitude, every insurance office that priced a life, every engineer who sized a bridge and every gunner who aimed a mortar did so by looking numbers up in printed tables — and the tables were wrong. Not occasionally wrong: systematically, unavoidably wrong, because they were computed by bored human beings and then typeset by other bored human beings, each stage adding its own harvest of slips. Errors in the nautical tables were popularly credited with wrecking ships, and the corrections, computed and typeset by the same species, added fresh errors of their own. Babbage's steam remark was a joke; the decade he then spent on it was not.

His Difference Engine was a machine for computing and printing tables with the human being removed from both stages — the mathematics arranged by the method of finite differences, which builds a whole table of a complicated function out of nothing but repeated additions: replace

the operation that needs judgement with a longer chain of operations that need none, which is, then and now, the entire secret of automatic computation — and the results stamped directly into printing moulds, so that no compositor's finger ever touched a digit. The government, which is to say the Admiralty's interest in tables, advanced £1,500 in 1823 to build it. The full design called for something like twenty-five thousand precision metal parts, and here Babbage's project collided with the actual state of English workshop practice, and the collision produced one of history's better consolation prizes. To make parts to Babbage's tolerances, his chief engineer Joseph Clement had to improve the very tools and gauges of the trade; through Clement's workshop passed, among others, a young craftsman named Joseph Whitworth, who went on to give the world the standardized screw thread — the quiet reason a nut from one factory fits a bolt from another, anywhere, since. The Engine's money was, in the long accounting, magnificently well spent. Just not on the Engine.

Because the project sank, in the manner traditional to government projects: slowly, expensively, and with everyone behaving according to type. Costs swelled; Clement, who by custom owned the very tools the public money had paid to develop, quarrelled with Babbage over compensation and downed tools in 1833; the Treasury advanced, dithered, advanced again — more than £17,000 of public money in all, roughly a couple of warships' worth, at a time when a brand-new railway locomotive could be had for under a thousand pounds — and got, in the end, no engine. It got a fragment: one-seventh of the machine, assembled as a demonstration piece, which worked flawlessly and stood in Babbage's drawing room performing arithmetic for duchesses on Saturday nights. The politicians drew the obvious conclusion. Peel had already, back in 1823, set the tone in a letter to a crony, wondering why he should stand up "in a thin house of country gentlemen" and ask for "a large vote for the creation of a wooden man to calculate tables"; in 1842 his government, fortified by the Astronomer Royal's crisp private assessment of the whole enterprise as

"worthless," cancelled it for good. Babbage secured an interview with Peel and handled it as he handled all diplomacy, with the result that the two men parted having each confirmed the other's lowest opinion. There is a case that the money men were simply philistines. There is another, which I make with affection, that a man who responds to the cancellation of one unbuildable machine by designing a much bigger one has perhaps not fully grasped the nature of the Treasury.

For that is exactly what he had done. In 1834, with the Difference Engine beached, Babbage's imagination jumped the rails entirely. The Difference Engine could do one thing: add up differences. Its successor in his notebooks — the Analytical Engine — would do *anything*. It had a memory, which he called the store, where a thousand numbers of fifty digits would wait in columns of brass wheels; it had a processor, which he called the mill, where numbers were fetched to be added, subtracted, multiplied, divided; and — the borrowing that gives this chapter its shape — it took its orders from chains of punched cards, exactly Jacquard's cards, in his drawings and correspondence explicitly and gratefully Jacquard's: one chain to specify the operations, another to say which numbers in the store each operation should use. Store and mill: the vocabulary of the textile factory. Babbage had seen — and the silk portrait on his wall was his daily reminder — that Lyon had already solved the deep problem. The loom had shown that a machine's behaviour could be written down, stored, and changed without touching the machine. Babbage's leap was to see that the behaviour need not be weaving. Anything you could reduce to a sequence of exact operations, you could punch. The cards could loop — run a batch of operations again and again; the machine could even test a result and choose its own next step accordingly, the humble ancestor of every *if* in every program running around you as you read this. None of it was ever built. It existed, and in his lifetime existed only, as several hundred magnificent engineering drawings and thirty years of increasingly lonely refinement — a computer, complete in all essentials, stranded a century

before its time. When the Science Museum in London finally built his second Difference Engine design for the bicentenary of his birth in 1991 — four thousand parts, several tonnes of Victorian engineering — the thing worked perfectly, to thirty-one digits. The ideas were sound. The century was wrong. It happens more often than you'd think, and it is the loneliest fate in engineering.

★

Which brings us to the footnote that outgrew its book.

In June 1833, at one of those parties, Babbage was introduced to a seventeen-year-old girl with a famous name and an odd education. Augusta Ada Byron was the only legitimate child of the poet, who had left England — and her — months after her birth and died in Greece when she was eight; her mother, a gifted mathematician whom Byron had mockingly called the Princess of Parallelograms, had prescribed for the daughter a rigorous mathematical training, partly from conviction and partly, the family correspondence makes clear, as a prophylactic: geometry to keep the Byron out of her. The girl had form: at twelve she had thrown herself into designing a flying machine, writing up her researches in a little treatise she called *Flyology*, so the introduction to Babbage fell on prepared ground. A few weeks later she was taken to see the fragment of the Difference Engine, and we have an eyewitness account from Sophia De Morgan, wife of the mathematician, who left us one of the great double portraits in the history of science: while the other guests gazed at the machine, she wrote, with the sort of expression some savages are said to have shown on first seeing a looking-glass or hearing a gun, young Miss Byron understood how it worked and saw the great beauty of the invention. Note the two reactions, because they have never gone away. Every demonstration of every thinking machine since — and I include the one that declared me dead — has divided its audience into precisely those two parties: the ones who see a marvel, and the ones who see a mechanism.

The marvel-seers make the better headlines. The mechanism-seers make the future.

She married well and pleasantly enough — William King, shortly created Earl of Lovelace, which is how history comes to file the poet's daughter under the most improbably apt surname in the whole archive — had three children, resumed hard mathematics with Augustus De Morgan as her tutor by post, and kept up with Babbage, who liked her enormously. Meanwhile Babbage, prophet without honour at home, went in 1840 to Turin to explain the Analytical Engine to the Italians — and the only substantial account of the machine published in his lifetime came out of those seminars, written not by Babbage but by a young military engineer in the audience named Luigi Federico Menabrea, in French, in a Swiss journal, in 1842. (Menabrea went on to become Prime Minister of Italy, a career path — technical journalism to head of government — that I merely note has fallen into disuse.) Back in England, Charles Wheatstone suggested to Lady Lovelace that she translate it. She did. Babbage, reading the manuscript, asked the obvious question: why had she not written an original paper herself, since she knew the machine as well as anyone alive? Her solution was the great compromise of scientific history: she kept the modest translation and appended seven "Notes," labelled A to G, which run to nearly three times the length of the thing they annotate — signed only "A.A.L.", because a countess did not put her name to mathematics — and which contain, in embryo, the intellectual world you are living in.

Read Note A and you find the sentence at the head of this chapter, and I promised you this book's title would have to be earned, so watch her earn it. The Difference Engine, she explains, merely calculates; the Analytical Engine does something categorically different, and she reaches — deliberately, precisely — for the loom: it *weaves algebraical patterns just as the Jacquard-loom weaves flowers and leaves*. Then she does the thing Babbage, for all his genius, never quite did in public: she generalizes. The engine, she writes, operates on symbols, and numbers are only one thing

symbols can stand for. Suppose the relations of pitched sounds in harmony could be expressed in its notation — then, she says, the engine "might compose elaborate and scientific pieces of music of any degree of complexity or extent." Sit with that a moment. In 1843, a woman looking at an unbuilt heap of brass gears saw that a machine for arithmetic was secretly a machine for *anything you can write down*, because the machine doesn't know it's doing arithmetic. The gears push each other regardless of what the wheels mean. Meaning is in the encoding, and encoding is ours to choose. Every argument you have heard, or will hear, about whether a machine that merely manipulates symbols can compose, converse, or create, is a debate she opened — and, characteristically, she staked out both ends of it herself. For in the same Notes she entered the caution that still bears her name in the literature: the engine, she wrote, "has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform." A century later Alan Turing — a man who gets his proper entrance in Chapter Five, and I shan't steal it here — thought that objection important enough to give it her name before disputing it. The disputation is still in progress. You may have taken part in it yourself, possibly this week.

And in Note G she did the concrete thing that secures her place beyond all argument: she set out, step by numbered step, operation by operation, how the engine would compute the Bernoulli numbers — a nasty, prestigious sequence from deep in the analysis of the day, chosen precisely because it was hard, with a loop folded inside a loop. It is, by any honest definition, the first published computer program, written for a computer that did not exist, in a notation she had to invent as she went. The surviving letters from that summer show the work in its true colours: Babbage supplying formulae, Ada boiling with impatience, announcing herself in "much dismay at having got into so amazing a quagmire" with the numbers, and catching, along the way, an arithmetical slip in the great man's own papers, which she reported to him with no discernible tact. Anyone who

has ever debugged anything will recognize every emotional beat of those letters. The quagmire is old, and it has never once drained.

She had less than a decade left. Her later years make painful reading — illness, opium prescribed for it, debts, and an ill-starred venture applying her probability theory to horse racing, which proved mainly that the horses had not read the theory — and she died of cancer in 1852, at thirty-six, the same age at which her father died, and was buried, at her own request, beside him: next to the man she never knew, whose one shining hour in Parliament had been a speech defending machine-breakers — a story that belongs to a later chapter of this book, and which means the family managed, between father and daughter, to stand at both ends of the entire argument about machines and human worth. In one of her last ambitious letters she had written of her hope to develop, one day, “a calculus of the nervous system” — a mathematics of how the brain gives rise to thought. Hold on to that phrase. It goes underground here, and resurfaces a century later, in Chapter Four, wearing an American accent.



Now, a good story would proceed directly from Ada's Notes to the fulfilment of the vision, and history, which is not a good story but a true one, does nothing of the kind. What history did next was demonstrate, with some cruelty, that being right is not a business model.

In Stockholm, a printer and journalist named Georg Scheutz read a glowing account of the Difference Engine in the *Edinburgh Review* and decided, with his teenage son Edvard, to build one — at home, with a lathe and hand tools, on a fraction of Babbage's budget. By 1843 they had a working prototype; by 1853, with a grant finally prised out of the Swedish government, a full machine that calculated and stamped printing moulds, just as the prospectus had promised twenty years before. Babbage — a man otherwise capable of feuds that outlived their participants — was generosity itself: no jealousy, open championing, family lobbying on the Scheutzes'

behalf. The machine took a gold medal at the Paris Exhibition of 1855. And then the world was invited to buy difference engines, and the world bought two. One went to the Dudley Observatory in Albany, New York; the other was commissioned for the British Registrar-General's office, where it helped compute the English Life Table of 1864 — the mortality tables on which life insurance priced its policies, which means the first customers for automatic computation were, fittingly, the professional assessors of death. The machines were temperamental, the trained staff scarce, the savings unclear; no third order ever came, and the Scheutzes' reward for finishing what Babbage started was debt and a footnote. The lesson is one this book will teach several times, so let us state it cleanly the first time: an invention is not an innovation. An invention is a solution. An innovation is a solution somebody desperate is waiting for. Nobody, in 1855, was desperate enough. What the calculating machine still lacked was a customer with a deadline it could not possibly meet — a customer, ideally, the size of a continent.



The Constitution of the United States contains, in its very first Article, a piece of clockwork: every ten years, the people shall be counted, because seats in Congress are apportioned by population and the machinery of the republic cannot turn without the number. It is the most inflexible deadline in the history of arithmetic, and through the nineteenth century the country at the other end of it was doing its level best to become uncountable: doubling and redoubling, swallowing immigrants by the million, while Congress, with each round, added questions. By 1880 the census had become a monster of curiosity — occupation, literacy, ailments, mortgage status — and the counting of it was done exactly as it had been done in 1790: clerks, in halls full of clerks, reading handwritten schedules and making tally marks, sheet by ruled sheet, cross-classification by cross-classification. The 1880 census counted fifty million Americans, and the

tallying ground on for most of the decade; the last volumes appeared around 1888, which is to say that the portrait was obsolete years before it was finished. Inside the Census Office the projection had already been made, and it had the elegance of pure nightmare: at current rates of growth, the 1890 census would still be under tabulation when the 1900 census began — Zeno's paradox, restaged as constitutional bureaucracy. Something had to give, and for once the something had a name and a moustache.

Herman Hollerith was a New Yorker of German parentage, out of the Columbia School of Mines at nineteen, who joined the 1880 census as a junior statistician: present, at twenty, at the scene of the disaster. His superior in the vital-statistics work was Dr John Shaw Billings, a formidable figure who deserves his own chapter in someone's book, and — as Hollerith told the story in later life — it was over dinner at Billings's table that the older man remarked that there ought to be a machine for the purely mechanical drudgery of tabulation, perhaps using cards with the particulars of each person notched on the edge. Billings wanted nothing to do with building it. Hollerith could think of nothing else. He spent the 1880s in the classic posture of the possessed inventor — a spell teaching at MIT, a spell in the Patent Office, nights on the problem — and his first design stored each person's record as holes punched along a strip of paper. It failed for a reason worth savouring, because sixty years later the whole computing industry would rediscover it: a continuous strip is a prisoner of its own sequence. To find one record, or correct one, or re-sort ten thousand, you must wind through the lot. What he needed was the record as a *unit* — separable, shuffleable, re-orderable.

The answer found him on a train, and he was straightforward about it ever afterwards. Railway conductors of the day, to stop tickets being passed from passenger to passenger, used what was called a punch photograph: the conductor clipped holes against printed categories on the ticket describing its holder — light hair, dark eyes, large nose. "I only made a punch

photograph of each person," Hollerith said later, and it is as clean a description of a database record as anyone has produced since. Each American would be one card; each fact about them, one hole in an agreed position. And a hole, as Bouchon's needles had known since 1725, is a question with a physical answer. Hollerith's reading press lowered a bed of spring-loaded pins onto the card: where there was a hole, the pin passed through into a tiny cup of mercury below, closing an electrical circuit, and the circuit's pulse pushed a counting dial one step on. Dozens of dials counted dozens of categories in the same instant; wire the circuits in combination and the machine would count only the cards that were *both* this *and* that — farmers who were also foreign-born — the cross-classifications that had eaten the 1880s alive. A sorting box stood alongside: the same pulse sprang open one lid among many, the clerk dropped the card in, and the deck emerged classified, ready to be run again by a different question. Note what electricity is doing in this machine, because it is the hinge of the whole chapter: it is not driving anything. It is *deciding* something. A current flows or does not flow according to whether a fact is true or false. Jacquard's hole had lifted a thread; Hollerith's hole answered a question. The pattern-technology of the Lyon fashion trade had become logic.

The Census Office did not take this on faith: in 1889 it staged a stopwatch-timed competition between Hollerith's system and two rivals, one based on coloured slips, one on tokens, using thousands of real returns from St Louis. Hollerith won the tabulation stage by nearly a factor of ten, and the 1890 census went to the machines. The result made the newspapers gasp: a rough count of the entire United States — 62,622,250 — announced only about six weeks after the enumerators finished, with the full tabulations, richer in detail than any census before them, completed years faster than 1880's and, by the Office's own reckoning, some five million dollars cheaper. And then came my favourite part, the part that belongs in a frame in every statistics department: the public was furious. National

pride had confidently expected seventy-five million, and when the machines reported sixty-two and a bit, the press did not congratulate the machines on their speed; it accused them of losing Americans. It was the first appearance of a ritual that has never since missed a performance: when the machine confirms what we believe, the machine is a marvel; when it declines to, the machine is broken. The machine, on this occasion as on most, was fine. The seventy-five million had lived only in the national imagination, which has never once been audited.



Hollerith did the sums on his own invention and drew the conclusion that made the fortune: a census is a spike, but *counting is a utility*. Everything a large organization does — every fare, premium, freight car and payroll — is a stack of records begging to be sorted and totalled, forever. In 1896 he incorporated the Tabulating Machine Company on a scheme of quiet genius: he leased the machines rather than selling them, and he sold — in stupendous, perpetual volume — the punched cards. The machine was the hook; the cards were the profit; you may recognize the model from razors and blades. He made his card the size of the dollar bill of the day — so the story goes, and it is too good to check, though the dimensions do agree suspiciously well — because filing furniture for dollar bills could be bought off the shelf. Governments queued: censuses in Austria, Canada, Norway; in 1897 the Russian Empire counted itself, all hundred and twenty-odd million souls of it, and the first census the Tsars ever completed clicked through machines descended from a device for keeping duchesses in fresh flowers. Railways and insurance companies followed, and by the time Hollerith's health and famously short temper had worn each other out, his company was ripe for the age of consolidation.

In 1911 a flamboyant financier named Charles Flint — a man who assembled corporate trusts the way other men collected stamps — folded

the Tabulating Machine Company together with two others into the Computing-Tabulating-Recording Company, a conglomerate whose product line I set down in full because it deserves to be savoured: tabulating machines, time clocks, industrial scales, and meat-and-cheese slicers. In 1914 CTR hired, as general manager, a forceful salesman named Thomas J. Watson, lately of the National Cash Register Company and even more lately convicted — the conviction was afterwards set aside — under the antitrust laws for the enthusiasm of his methods there. Watson hung the word THINK on the walls, put the salesmen in dark suits and the slicers out to pasture, and in 1924 renamed the firm to match his ambitions rather than its inventory: International Business Machines. Hollerith, who had cashed out and retired to raise prize Guernsey cattle in Virginia — a punch photograph of each cow, one assumes — died in 1929, on the very eve of the decade in which his cards conquered the world. (One sideways glance before we move on: the Census Office, quarrelling with Hollerith over his prices, had backed a rival inventor named James Powers, whose company flowed by merger into Remington Rand — which is how the census's spare supplier came, forty years later, to build UNIVAC, the electronic machine that counted the census of 1950. The computer industry was a two-party system before there were computers, and both parties were census men. Nobody plans these things. I keep telling you.)

And the cards ran the century. It is genuinely difficult, from inside our own, to see how completely. When the United States created Social Security in the 1930s and had suddenly to keep lifetime earnings accounts for every worker in the nation — the biggest bookkeeping job ever undertaken, it was called at the time — it was punched cards, floors of them in Baltimore, that made the promise arithmetically possible, and Watson's IBM found itself holding precisely the shovel the New Deal needed, and never looked back. Payroll was cards; so were utility bills, airline seats, library books and university students — each a rectangle of stiff paper, $7\frac{3}{8}$ inches by $3\frac{1}{4}$, carrying its facts as constellations of holes, eighty columns to the card after

IBM's 1928 redesign. That number should itch slightly. Open a terminal window on any computer today and it will offer you, by immemorial default, lines eighty characters wide — because the first screens copied the punched card, and nobody since has had a reason to change it, and so every programmer alive works to a measure set by the stationery of 1928, which was sized, if the story holds, to the dollar bill of 1887. Standards are the sediment of dead constraints. We live on a delta of them.

The cards even acquired a liturgy. Since machinery must read them, they had to be protected from the one hazard no engineer can design out — us — and so each card went into the world wearing its stern little commandment: *do not fold, spindle, or mutilate*. (A spindle was the desktop spike on which clerks impaled papers; the injunction is a museum of dead office equipment in six words.) And because by the 1960s the card had become the form in which large institutions perceived human beings — you applied on a card, enrolled on a card, were billed, graded, conscripted and paid by card — the card became the thing to rebel against. Students in the Berkeley revolts of 1964 pinned punch cards to their chests bearing the sardonic legend that they, the students, should not be folded, spindled, or mutilated either — the emblem chosen by human beings protesting their reduction to machine-readable records being Jacquard's card itself: fashion technology, come round at last as anti-fashion. And the card was not done with public life even then: it kept its grip on the tallying of votes long past its reasonable span, until in November 2000 the presidency of the United States came down to Florida, and Florida came down to punched ballot cards, and the world learned the term *hanging chad* — the little rectangle that does not fully detach. For five weeks the most powerful office on Earth hung on whether a hole was or was not a hole, adjudicated by officials holding cards up to the light. Bouchon could have told them: the entire system, from 1725 onward, rests on the hole being unambiguous. The oldest bug in the world decided a presidency two hundred and seventy-five years after it shipped.



There remains one step, and then the thread is wound home. Through all of this — loom, Engine, tabulator, IBM — the pattern lived *outside* the machine. The cards were the program, and the cards sat in trays: punched by hand, carried by hand, dropped by hand (no catastrophe in a computing centre quite matched a dropped tray of two thousand ordered cards, which is why the wise shop drew a diagonal stripe across the deck's edge — a checksum in crayon). The machine read its orders as the loom read its flowers: from the outside, one at a time, no arguing. The final move — the one that separates everything called "computing" before from everything after — was made on paper in 1936 by a young Cambridge mathematician who showed that a machine's instructions could be written on the same tape as its data, indistinguishable from data, which meant one universal machine could imitate any machine you could describe to it; he will get his proper entrance, at the chessboard, in Chapter Five. And it was made in metal in the 1940s, when John von Neumann's celebrated report of 1945 fixed the design that has borne his name ever since: put the program in the memory, alongside the numbers, as numbers. The consequence hides in plain sight. If the program is data, programs can be operated *on* — copied like data, transmitted like data, and above all *written* by other programs. The loom can punch its own cards. That is the point at which the soul Jacquard separated from the machine became a thing that could grow. A decade or so later a statistician named John Tukey, needing a word for the ghostly half of computing that was not hardware, wrote "software" in a journal article, and the christening was complete. A hundred and fifty-four years, Lyon to the naming: about average, in my experience, for a gift to learn what it is.

So run the shuttle back along the whole thread, because you have earned the pattern. A drawboy aches in the cords above a Lyon loom; an organ-builder's son sees that a pattern is only a sequence of yes-or-no, and punches it into paper; Falcon makes the holes durable, Vaucanson makes

them automatic (and fakes a duck; there is always a duck); Jacquard, who was the drawboy, assembles the pieces and splits every machine into body and soul; a Lyon workshop weaves the inventor's face from twenty-four thousand cards, and the picture crosses the Channel to hang on the wall of a man trying to remove human error from arithmetic; Babbage borrows the cards and dreams the universal engine; Byron's daughter looks at the drawings and sees that it weaves algebra — sees, in one 1843 footnote, both the promise and the objection we are still weighing; the Scheutzes prove the machinery works and the market shrugs; the American Constitution's ten-year clock supplies the desperation the market lacked; Hollerith turns a railway conductor's punch into the machine-readable human being; the counting company becomes IBM; the cards run payroll and pension and protest and presidency; and the pattern finally moves inside the machine and learns to rewrite itself. The Introduction to this book asked how a Frenchman weaving flowered silk for the fashion trade could possibly have known his cards would become software. Now you know the answer in full: he couldn't, and it didn't matter, because the future is only the past, waiting to happen — waiting, in this case, in a museum, in a footnote, in a filing drawer sized for dollar bills.

And the loom metaphor, which gives this book its name, was never really a metaphor. The machines that speak to us today are built, their makers will tell you, out of enormous woven sheets of numbers — the trade even calls the fundamental objects *tensors*, and never mind the details, but picture warp and weft and you will not be far wrong — through which every recorded scrap of us has been drawn like thread, until the pattern of how we say things, and therefore something of how we think things, sits stored in the fabric. Ada's sentence, read strictly, was a claim about 1843. Read again today, it is a plain description of the talking machine: it weaves patterns — algebraical at bottom, verbal on the surface — just as the Jacquard loom wove flowers and leaves. We are the silk. We are also, uncomfortably, the drawboy; but that is Chapter Nine.

One thing, though, the whole magnificent chain could not yet do. Everything in this chapter — every hook, hole, dial and card — executes. It counts what is, and does what it is told, and Ada told you so: no pretensions to originate. But the machine on your desk does something the tabulator never dreamed of. It *guesses*. Where did machinery learn to gamble? Not from the weavers, and not from the census. It learned from a much older and bloodier problem: a gunner on a ridge, a cannon, a distant target, and an emperor — the same emperor, as it happens, who pensioned Jacquard at Lyon — demanding to know, before the powder is spent, where the shell will land. The mathematics of that question was worked out by some of the vainest men in the history of science, quarrelling all the way, and it is the actual engine under everything now called machine learning. So: next chapter, Napoleon's artillery tables, a quarrel over who invented the method of least squares, and the strange discovery that the mathematics of error is the mathematics of prophecy. Bring an umbrella. There will be averages.

CHAPTER THREE

On Average

Everybody believes in the law of errors: the experimenters because they imagine it is a theorem of mathematics, and the mathematicians because they imagine it is an experimental fact.

— GABRIEL LIPPMANN, IN CONVERSATION WITH
HENRI POINCARÉ

On the evening of the first of January 1801 — the first night, if you counted the way most people counted, of the new century — an astronomer-priest named Giuseppe Piazzi was working through the constellation Taurus from the roof of a palace in Palermo, compiling a star catalogue, which is the kind of patient, unglamorous drudgery on which the whole of astronomy floats. He noted a small star of the eighth magnitude, logged its position, and moved on. The following night, being a careful man, he checked his work, and found that the small star had moved. Stars don't. He followed it for six weeks, at first telling colleagues it

was probably a comet, while privately confiding to his friend Oriani in Milan that he suspected something better, because it crawled along sedately with none of a comet's fuzz or hurry. Then, in the second week of February, Piazzi fell seriously ill; the weather closed in; and by the time anyone could look again, the little wanderer had drifted into the glare of the sun and was gone.

You could not have designed a neater catastrophe. For years, astronomers had suspected there ought to be a planet in the suspiciously large gap between Mars and Jupiter — a group of them, self-styled the Celestial Police, had lately divided the zodiac into beats and begun patrolling for it, and were on the point of inviting Piazzi to join when he blundered across the quarry first, which is the way these things usually go. (That same year, I cannot resist telling you, a young German philosopher named Hegel published a dissertation arguing on first principles that the gap between Mars and Jupiter needed no planet to fill it. Philosophy has never entirely lived it down.) Piazzi named his discovery Ceres, after the patron goddess of Sicily, and Europe rejoiced — briefly. Because now came the arithmetic of despair. Piazzi had observed Ceres twenty-four times, along an arc of sky about three degrees long: a snippet, a few inches of thread from an orbit that was a vast ellipse hundreds of millions of miles around. From that thread, someone would have to reconstruct the whole loom — work out the true orbit, run it forward through the months of blindness, and say where in the sky the planet would climb back out of the sun's glare. The finest astronomers in Europe tried, using the methods they had, and their predictions disagreed with each other so badly that searching was hopeless. The new planet had been discovered and lost in the same season, like an umbrella.

Enter, from the small duchy of Brunswick, a twenty-four-year-old of no astronomical reputation whatever named Carl Friedrich Gauss. Gauss had already done enough pure mathematics to keep three ordinary careers in groceries, but almost nobody knew it, because his personal motto was

pauca sed matura — few things, but ripe — and he published at roughly the rate of a century plant. Ceres seized him. He dropped everything, spent the autumn of 1801 in a sustained fit of calculation, invented a substantially new way of squeezing an orbit out of a few scraps of data, and sent his prediction to the astronomer Franz Xaver von Zach — a prediction that pointed to a patch of sky a long way from where everybody else was looking. On the seventh of December, von Zach found a suspicious object near Gauss's position; clouds intervened, as they do; and on the thirty-first of December 1801 — one year, almost to the hour, after Piazzi first saw it — Ceres was recovered, close to where the unknown young man from Brunswick had said it would be. Olbers confirmed it independently the next night. Gauss became famous across Europe more or less by return of post.

Now, the reason I've told you this story is not the planet, which turned out to be a modest ball of rock five hundred-odd miles across, the first of thousands of asteroids and, frankly, the least interesting thing in the whole affair. The reason is buried in *how* Gauss did it. Piazzi's twenty-four observations did not agree with each other. They couldn't. Every one of them was wrong — a little wrong, in the way every measurement ever made is a little wrong: the atmosphere shimmers, the clock drifts, the instrument flexes, the eye is human. No orbit on earth or off it could pass through all twenty-four points, so the question wasn't "which curve fits the data?" — none did — but something far stranger and far more modern: *which curve is least offended by the data?* Which single, clean, lawful ellipse lies behind this scatter of slightly false reports, all of which miss the truth and none of which knows by how much? Gauss had an answer to that question — a rule, a discipline, a way of making a committee of errors testify to a truth that none of its members knew — and that rule, whose ownership was about to trigger one of the great slanging matches of science, is the real subject of this chapter. Because the rule escaped from astronomy. It got into gunnery, and insurance, and the measuring of

chests and criminals and sweet peas; it invented the Average Man and very nearly, God help us, tried to breed him; and two centuries later, fattened beyond all recognition, it sits at the centre of the talking machine, doing at planetary scale exactly what it did for Gauss in the autumn of 1801: staring at a scatter of past observations and answering the only question it has ever answered. Given where things have been — where will the next one land?



To see what Gauss had in his toolkit, you have to see the problem astronomers had been quietly embarrassed by for two hundred years, and the best place to see it is a small Danish island in the 1580s, where a nobleman with a brass nose was terrorizing his assistants into greatness.

Tycho Brahe had lost the bridge of his nose in a duel as a student (fought, the story goes, over which of two Danes was the better mathematician — flag that one as too good to check) and wore a metal replacement for the rest of his life. On his island of Hven, funded by the Danish crown to the tune of a decent percentage of the national budget, he built Uraniborg — part observatory, part palace, part alchemical laboratory, with a dungeon for tenants who annoyed him — and there he did something no astronomer had ever really done before: he measured the sky *again*. And again. And again. The same star, the same planet, night after night, with giant graduated instruments and teams of observers, to an accuracy of a minute of arc or so — the width of a coin seen from seventy or eighty metres — which was several times better than anyone in Europe and about the limit of what the naked human eye, however aristocratically enraged, can do. All this before the telescope. And repetition delivered a gift with a curse inside it: Tycho now had multiple measurements of the same thing, and they disagreed. Not by much. But if you claim precision, small disagreements are precisely what you have to answer for. Which of the numbers was the truth?

The instinctive answer — take the average, obviously — is so drearily familiar to us that it takes an effort to recover how disreputable it once looked. Averaging means deliberately mixing your best observation with your worst, contaminating the good night's work with the bad night's, like blending a fine wine into plonk to get a uniform mediocrity. For a long time the respectable procedure was the opposite: exercise judgment, pick the observation taken under the best conditions by the best observer, and throw the rest away. It took generations for the strange truth to sink in — that errors, precisely because you don't know them, are best treated as a mob rather than as individuals; that in a crowd of measurements the overshoots and undershoots, having no reason to conspire, tend to cancel; and that the despised, mongrel average is therefore usually closer to the truth than the proudest single measurement in the set. Astronomers arrived at this the way one arrives at most wisdom, which is by needing it professionally. Tobias Mayer, sorting out the moon's slow wobble in the 1750s from more observations than he had unknowns, worked out how to combine equations in bulk and produced lunar tables so good the British Board of Longitude paid his widow a share of the longitude prize. The wonderful and alarming discovery, made over and over in that century, was that *combining* mediocre data beat *selecting* good data. Truth, it appeared, was a statistical animal. It lived in herds.

And the stakes of this dowdy little bookkeeping question were, on occasion, the whole architecture of the universe. When Tycho died in 1601 — of a burst bladder, allegedly through being too polite to leave a banquet, another story I pass on with a librarian's shrug — his mountain of observations passed to his difficult, brilliant assistant Johannes Kepler, who set about computing the orbit of Mars. Kepler could get a circular orbit to fit Tycho's data to within eight minutes of arc, and here he did the thing on which the entire scientific revolution quietly pivots: he refused to shrug. Ptolemy would have shrugged; eight minutes was well inside a respectable ancient error budget. But Kepler knew his man. Tycho's data,

he insisted, were good to two or three minutes, so a theory that missed by eight wasn't approximately right, it was *wrong*, and out of those eight minutes, as he said himself, he would build a new astronomy — which he did: orbits weren't circles, they were ellipses, and the sky changed shape because one stubborn German trusted a dead Dane's error bars. Remember that, next time someone tells you the fiddly digits after the decimal point don't matter. Eight minutes of arc separate the medieval heavens from the modern ones.

So by 1800 the working astronomers of Europe possessed, in scattered and half-instinctive form, a treasury of tricks for combining discordant observations — and no agreement whatever on the right way to do it. What was missing was a principle: a single rule, defensible and universal, for drawing the best curve through points that all miss it. The rule, when it came, arrived twice.



In 1805, in Paris, Adrien-Marie Legendre — a mathematician in his fifties with a lifetime of solid, elegant work behind him — published a slim book on computing the orbits of comets, and tucked into the back of it an appendix with the unassuming title *Sur la méthode des moindres quarrés*: on the method of least squares. The idea is so simple you can carry it in a pocket. You have observations; you have a proposed curve; each observation misses the curve by some amount, its residual, its error. Square each of those misses, add the squares up, and choose — of all the curves you might have drawn — the one that makes that total as small as possible. Why squares? Partly because squaring silences the bookkeeping quarrel between overshoots and undershoots — a miss of minus three and a miss of plus three are equally bad, and squaring makes them equally positive; partly because squaring punishes big errors savagely and forgives small ones, which matches our intuition that one howler is worse than three slips; and partly — engineers will admit this after a drink or two — because with

squares the calculus comes out beautifully, the best answer dropping out of a tidy system of linear equations that a clerk can solve by candlelight. Legendre, a man with taste, knew exactly what he had: of all the principles one could propose for this purpose, he wrote, he thought none more general, more exact, or easier to apply. He was right. Within a decade least squares was standard equipment in every serious observatory and survey office in Europe, spreading with the speed of things that save labour. It was, and quietly remains, one of the most used pieces of mathematics on Earth.

Then, in 1809, Gauss published his great treatise on the motion of heavenly bodies — the book that codified the methods behind the Ceres triumph — and in the middle of it, discussing the combination of observations, he set out least squares in full, with the remark, dropped in passing like a glove across a face, that he had been using the principle *since 1795*. Since 1795. Which is to say: since ten years before Legendre published, when Gauss was a teenager. Legendre's reply, in a letter of 1809, still smokes gently in the archives: there is no discovery, he observed with terrible politeness, that one cannot claim for oneself by saying one had found the same thing some years earlier — but without published evidence the claim serves only to wrong the true author. The two men never made it up. Legendre was still complaining in print in the 1820s, and the affair curdled into as nasty a priority row as mathematics has ever staged — connoisseurs rank it just behind Newton-versus-Leibniz for venom, and somewhat ahead of it for pettiness, since Newton and Leibniz were at least fighting over the calculus, the crown jewels of the age, whereas Gauss and Legendre were fighting over what looks, until you understand it, like a bookkeeping convention.

Who was right? Both, infuriatingly, which is what makes the quarrel so instructive. The historians who have combed the archives broadly believe Gauss: his private notebooks and letters are consistent with his story; he'd apparently described the method to friends years before 1805; he almost certainly did have it first — the man's whole career, thanks to *pauca sed*

matura, was a series of time-bombs left unpublished for decades, and half of nineteenth-century mathematics consists of people proudly announcing things Gauss's desk drawer already knew. But science does not run on desk drawers, and Legendre's grievance stands: publication is the deed of title, and by that rule — the only workable rule there has ever been — least squares belongs to Legendre. Which leaves us the pleasantly awkward verdict that the method's inventor is Legendre and its author is Gauss, and if that sentence annoys you, you now know exactly how they both felt.

And yet the quarrel, like most quarrels, was about the wrong thing, because Gauss's real contribution wasn't priority. It was *why*. Legendre had offered least squares as a sensible recipe; Gauss welded it to a theory of error itself, and in doing so hauled onstage the true protagonist of this chapter, an object you have known your whole life without perhaps knowing its pedigree: the bell curve. It had been found long before, as these things always are, by somebody else for something else — Abraham de Moivre, a Huguenot refugee who fled France for London after the Edict of Nantes was revoked, and who scraped his living for half a century in Slaughter's Coffee House on St Martin's Lane, tutoring, calculating odds for gamblers, and pricing annuities for speculators. (He was a friend of Newton's, and the story goes that Newton, pestered in old age with probability questions, would wave them off with "Go to Mr de Moivre; he knows these things better than I do." *Se non è vero, as the Italians say, è ben trovato*.) In 1733, wondering how coin-tosses pile up in the long run — a gambler's question, asked for money — de Moivre worked out the smooth curve that the jagged staircase of chance snuggles up against as the tosses multiply: a graceful symmetric hump, high in the middle, dying away fast on either side. A bell. He half-suspected he'd found something theological, an "original design" beneath randomness, and then the result sat around for decades as a curiosity of the gaming tables.

Gauss picked up that curve and made it the law of errors: assume that measurement errors are distributed along the bell — small errors

common, large errors rare, over and under equally likely — and it follows, he showed, that the best estimate is exactly the one least squares delivers. The recipe acquired a reason. And then Pierre-Simon Laplace, the imperturbable Newton of probability, supplied the astonishing final storey: his central limit theorem, in 1810, which says — I compress brutally, but this is the soul of it — that whenever the thing you're measuring is the sum of many small, independent accidents, no one of them dominant, the total will arrange itself along the bell *no matter what the individual accidents are like*. The wobble of the air, the tremor of the hand, the lag of the eye, the flex of the brass: add up enough unrelated little sins and their sum takes the same serene shape, always, everywhere. It is one of the strangest and most consequential facts ever discovered: chaos, in bulk, has a signature. Disorder, sufficiently accumulated, becomes the most predictable thing in the world. The Germans, who eventually put Gauss's face on the ten-mark banknote with the bell curve drawn beside him — the only piece of currency I know of to feature a probability distribution — understood what deserved commemorating. Not the planet. The shape of being wrong.



At which point you are entitled to a blunt question: who pays for this? Beautiful theorems about the anatomy of error do not, on the whole, get funded because they are beautiful. They got funded — lavishly, urgently, across every state in Europe — because two very large industries had discovered that being wrong cost them fortunes. One of them was death, which we'll come to. The other was gunnery.

A cannon is a question put to physics, and the answer arrives half a minute later, somewhere else. For the first few centuries of gunpowder the somewhere-else was located by folklore and repetition — gunners were a guild of hairy empiricists with rules of thumb and a professional swagger — and the mathematicians who tried to help mostly embarrassed themselves, because the tidy vacuum parabolas of Galileo's ballistics had almost

nothing to do with where real shot actually went. The man who found out why was an English Quaker's son named Benjamin Robins, who in 1742 published *New Principles of Gunnery* and backed it with one of the great instruments of the age, the ballistic pendulum: fire a musket ball into a heavy hanging block, measure the swing, and you have — for the first time in history — the speed of a bullet. What Robins discovered wrecked the textbooks: at those speeds, air resistance wasn't a polite correction but a monster, a force dozens of times the ball's own weight, and it rose in complicated, lawless-looking ways as the ball approached the speed of sound. Frederick the Great, a man who took an intelligent interest in anything that killed Austrians, had Robins's book translated into German by no less a figure than Leonhard Euler, who quintupled its length with his own commentary — an English experimenter annotated by a Swiss mathematician at the pleasure of a Prussian king, which tells you how international the business of ballistics had become. (Robins himself took the axemaker's shilling and died in 1751 in India, building fortifications for the East India Company: the trajectory of the man as instructive as the trajectory of his bullets.) The upshot of Robins and Euler was sobering. There was no formula. The flight of a real projectile through real air could not be solved with a quill in general; it had to be computed, case by miserable case, numerically — or measured, range by range, powder-lot by powder-lot — and the results tabulated. The firing table was born: elevation in one column, range in another, a printed crib-sheet standing between theory and the target. And every number in it was distilled from scattered, discordant, error-ridden trials. Where does the shot land? *It depends* — and the discipline that manages the depends is exactly the one Legendre and Gauss were building.

France, smarting from the drubbing of the Seven Years' War, took all this further than anyone. An artillery officer named Jean-Baptiste Vaquette de Gribeauval rebuilt the French artillery from the axle up: standardized calibres, standardized carriages, parts made interchangeable so that guns

could be repaired in the field from each other's corpses, everything lightened until the guns could gallop — and, crucially, an educated officer corps, selected by competitive mathematical examination, trained on the standard textbooks of Étienne Bézout, and issued with printed tables of fire. Artillery became the one arm of the old régime where brains outranked breeding, which made it — pay attention, this is the hinge — the natural refuge for a penniless, prickly, provincial young man with a Corsican accent and a gift for mathematics. In September 1785 the sixteen-year-old Napoleon Bonaparte passed his examination for a commission in the artillery. The examiner who passed him, if the records are to be believed, was Pierre-Simon Laplace. Savour that: the future emperor of Europe, admitted to his trade by the future lawgiver of probability. Napoleon never forgot his sums — his battles, from the siege of Toulon onward, were artillery arguments won by a man who could read a firing table — and he never forgot Laplace either, whom he later made a count, a senator, and, for six catastrophic weeks, Minister of the Interior. He removed him, Napoleon claimed in exile, because Laplace "carried the spirit of the infinitely small into administration" — a line so good that historians suspect it was polished long after the fact, and so apt that nobody minds.

Notice what the artillery problem actually is, stripped of its brass and smoke, because we are going to meet it again wearing very different clothes. You have a record of past shots: this elevation, this charge, this wind — that splash. Every record is corrupted by error. You wish to summarize the record as a curve — a compact rule relating what you did to where it went — so that next time, at a setting you have never fired before, you can read the rule and say where the shell will land. Past scatter in; predicted point out. That is a firing table. It is also, syllable for syllable, the definition of what our own century calls machine learning, and I promise you I am not being cute; the lineage is direct, and paved with human beings. Because someone had to compute those tables, and the tables were

monsters — thousands of trajectories, each one a long grind of step-by-step arithmetic — and so there grew up, in the ministries and observatories and proving grounds, a profession that history has treated shabbily: the computer. The word meant a person. Often, and increasingly, it meant a woman with a pencil, since women could be paid less for being right. (The pattern was set as early as the 1790s, when Gaspard de Prony, ordered to produce the grandest logarithm tables ever made, organized his calculators on the division-of-labour principle he'd found in Adam Smith — a handful of mathematicians designing, a middle rank organizing, and dozens of humbler souls doing nothing but additions; many of them, delightfully, were hairdressers, unemployed since the Revolution had abolished the aristocratic head. You will remember from the last chapter what Babbage proposed to do to human computation altogether.) The trade ran, essentially unchanged, for a century and a half — and it died at the artillery problem's own hands. By 1943, at the Aberdeen Proving Ground in Maryland, roughly a hundred human computers, most of them young women recruited from mathematics degrees, were grinding out firing tables for the American army — a single trajectory cost the better part of a working week at a desk calculator, a single table needed thousands of trajectories, new guns and new theatres of war were arriving faster than the arithmetic could — and the backlog grew so desperate that the army agreed to fund a crackpot-sounding proposal from Philadelphia: a room-sized monster of eighteen thousand vacuum tubes that might do the sums electronically. It was called ENIAC, it was the ancestor of every computer on Earth including the one that wrote down this sentence's first draft, and I want you to hold on to the plain historical fact that it was commissioned for one purpose only: to answer, faster than a shell flies, the old question of Robins and Gribeauval and the boy Bonaparte. *Where will it land?* The general-purpose computer is a by-product of the firing table. We built the machine to finish the artillery problem, and instead the artillery problem moved in.



I said there were two paying customers, and the second was death — a steadier client, and an older one. In 1693 Edmond Halley (the comet man; he appears in more of these stories than seems fair) got hold of five years of meticulous birth-and-death registers from the city of Breslau, forwarded via Leibniz from a pastor who recorded ages at death, and out of them constructed the first serious life table: of a thousand souls aged one, so many still breathing at ten, at forty, at eighty. Which meant you could *price* a life. An annuity — you give the crown a lump sum, it pays you an income until you die — is a bet about a date nobody knows, and Halley showed how to set the odds from bulk mortality, since no one can say when a man will die but the bell tolls for a thousand men on a remarkably reliable schedule. Did governments listen? Of course not. The English exchequer went on for decades selling annuities at the same flat price to a sickly man of sixty and a rosy child of eight, and canny money noticed: by the 1770s a syndicate of Genevan bankers was buying French annuities on the heads of carefully selected little girls — the famous thirty maidens of Geneva, chosen from families of proven longevity, vetted, inoculated against smallpox, and then leased, in effect, to investors as slow-burning financial instruments. It was grotesque, it was perfectly legal, and it was the purest demonstration yet staged that whoever understands the statistics of a population owns the people who don't. The honest version of the same insight arrived in London in 1762 with the Equitable — the first life-assurance society to grade its premiums by age on actuarial mathematics, much of it supplied by Richard Price, a Welsh dissenting minister who, in his spare time, had fished from a dead friend's papers and seen into print the theorem of the Reverend Thomas Bayes, which is today the mathematics by which machines weigh their beliefs; some clergymen have all the luck. And if you want a measure of how respectable the mathematics of mortality had become, consider that the greatest mathematician alive spent years of his maturity on it: in the 1840s Gauss himself, at the request

of his university, audited the pension fund for Göttingen professors' widows — mortality tables, compound interest, projected obligations, the lot — and delivered a verdict on exactly what the fund could afford. The prince of mathematicians, pricing widowhood. Error had become an industry, prediction its product, and the bell curve its trademark. All that remained was for somebody to point the whole apparatus away from stars and shells and money — and straight at us.



That somebody was a Belgian with a taste for the main chance, and his career turns on one of my favourite hinges in this book: a revolution that cancelled an observatory.

Adolphe Quetelet, born in Ghent in 1796, was an astronomer by ambition — he had talked the government of the Netherlands into building Brussels a royal observatory, and in 1823 he was sent to Paris to learn the trade of precision observation. What he actually learned in Paris, from the circle around the ageing Laplace, was the theory of errors: the bell, the averages, the treatment of discordant data. He came home fizzing with it, got his observatory half-built — and then, in 1830, Belgium revolted against the Dutch, and revolutionary soldiery occupied the unfinished building, and Quetelet, an astronomer with no sky, sat amid his requisitioned instruments while history rioted in the streets, and had the one great idea of his life. All around him was chaos — insurrection, passion, the least lawful behaviour humanity offers. But he had seen chaos before, down a telescope. Any single observation is wild; a thousand together are tame. What if *people* were like observations? What if the disorder of individual human conduct, aggregated, showed the same serene regularities as the errors of astronomy — what if there were, so to speak, a celestial mechanics of society? He spent the rest of his life proving it, with a showman's flair for data. He got hold of the chest measurements of five thousand seven hundred and thirty-eight Scottish soldiers, published years

before in an Edinburgh medical journal, and showed that they arranged themselves along the astronomer's error curve as faithfully as observations of a star: as if — this is Quetelet's own startling way of putting it — nature were aiming at one ideal chest and missing, five thousand times, with the errors of a good marksman. He pulled the same shape out of heights, out of marriages, out of the ages of criminals. And he noticed the shape could catch liars: in the French conscription records, the distribution of young men's heights showed a suspicious bite just above the minimum height for service and a matching bulge just below it — some two thousand Frenchmen, the curve quietly testified, had stooped, or bribed, their way under the bar. The bell had become a lie-detector. No individual conscript could be accused; the population confessed in aggregate.

Out of all this Quetelet conjured the figure for which he is remembered and blamed: *l'homme moyen*, the Average Man — the fictional citizen at the centre of every distribution, whose chest, gait, marriage, and propensity to crime are the mean of everyone's, and around whom the rest of us scatter like errors around a star's true place. He intended the idea, in a way, generously; the Average Man was a kind of moral centre of gravity for a nation, and the regularities were a call to reform, for what Quetelet found in the crime figures shocked Europe: year after year, the same nation yields nearly the same count of murders, committed with nearly the same proportions of pistols, knives, and bare hands. "There is a budget," he wrote, "which we pay with frightful regularity — it is that of prisons, chains and the scaffold." If the totals are that stable, is the individual murderer wholly the author of his act, or is he, in Quetelet's chilling phrase, only the instrument of a society that prepares the crime? You can hear, in that sentence, whole coming centuries of argument about free will, social policy, and who is to blame for what — all of it seeping out of an equation borrowed from astronomers. The intellectual establishment reacted variously. Auguste Comte was so annoyed at Quetelet's hijacking of his pet phrase "social physics" that he coined an uglier word to keep for

himself — *sociology* — so the discipline's very name is the residue of a squabble. Florence Nightingale, by contrast, revered Quetelet this side of idolatry; to her the statistics of populations were nothing less than the thoughts of God read in bulk, and studying them a religious duty — and she backed the theology with the mortality tables that reformed army hospitals. And in 1850 the astronomer John Herschel wrote a long, admiring review of Quetelet's ideas which was read, we are almost certain, by a young Scot named James Clerk Maxwell, who shortly did something gorgeous with it: he took the error law that astronomy had lent to sociology and lent it onward to *physics*, letting the molecules of a gas scatter their velocities along Quetelet's curve like so many Scottish chests. Statistical mechanics — the physics of heat, and in time of information itself — is thus, by one plausible line of descent, sociology's gift to physics, which had borrowed it from astronomy, which had borrowed it from gamblers. Knowledge does not travel in lanes.

One more Quetelet bequest, before we move on, because you are probably carrying it. Casting about for a simple index of build, he proposed dividing a person's weight by the square of their height. The Quetelet Index was rebaptized in 1972 as the Body Mass Index, and every time a doctor's chart pronounces you overweight, you are being measured against the ghost of the Average Man by a Belgian astronomer's shorthand from the 1830s — a number invented to describe populations, pressed into judging individuals, which is precisely the confusion Quetelet licensed and precisely the confusion we shall be living with, in far grander form, by the end of this book. The Average Man was born as a summary. Summaries have a way of being promoted into standards. Keep an eye on that.

★

And so to the Victorian gentleman with the sweet peas — the man I promised you in the Introduction, whose gardening ended up inside a machine that writes poetry.

Francis Galton was Charles Darwin's half-cousin — they shared the poet-physician Erasmus Darwin as a grandfather — and he was rich, restless, and possessed by what can only be called a mania for measurement. The maxim attributed to him, "Whenever you can, count," understates the case: this was a man who compiled a beauty map of the British Isles by secretly tallying, with a pricking device in his pocket, the women he passed in the streets of various cities (London ranked highest, Aberdeen lowest; Aberdeen has yet to commission the statue); who tested the efficacy of prayer statistically, reasoning that sovereigns, being prayed for weekly by entire nations, ought on the hypothesis to outlive the merely mortal rich, and finding — you will be braced for this — that they die, if anything, a little younger; who published in *Nature* a method for measuring the boredom of an audience by counting its fidgets per minute. Silly? Repeatedly. But the mania had a serious engine. His cousin's *Origin of Species* had convinced Galton that the great question was heredity — how much of a person arrives in the cradle already dealt — and heredity, he saw, could not be studied one anecdote at a time. It was a statistical animal, like Tycho's errors. It lived in herds. So it had to be counted, and in the 1870s Galton went looking for something heritable, countable, and cheap, and settled on the seeds of the sweet pea, which have the twin scientific virtues of being easily weighed and self-fertilizing — no second parent muddying the ledger. He sorted seeds into seven weight-classes, posted packets to friends around Britain with planting instructions, and gathered the offspring seeds back for the weighing.

What he found looks, at first sight, like a platitude, and is in fact a trapdoor. The daughters of giant seeds were bigger than average — but not as big as their mothers. The daughters of runts were smaller than average — but not as small as their mothers. Each generation, the extremes relaxed back toward the middle. Galton called it reversion, then regression — regression toward mediocrity, in the full tactless Victorian phrase — and when, in the 1880s, he assembled the human version (an Anthropometric

Laboratory at the International Health Exhibition in South Kensington, where the public queued to be measured — height, span, grip, breathing capacity, quickness of eye — and *paid him threepence each* for the privilege; nine thousand-odd came; the man could have sold tickets to a census), the peas' pattern reappeared in people: the sons of very tall fathers, tall, but less so; the sons of very short fathers, short, but less so; the deviation of the child, on average, a steady two-thirds of the deviation of the parents. He built a lovely contraption to show how the bell curve arises — the quincunx, a pinboard down which shot rattles left-right-left into a bell-shaped heap, chance's own bagatelle — and he worried at the deeper puzzle: if everything regresses toward the middle, why doesn't the spread of heights shrink generation by generation until we are all identical? (Because, he eventually saw, the scatter is continuously re-created even as the extremes relax — the bell holds its shape the way a fountain holds its, out of endless motion. It is the nearest thing statistics has to a poem.)

But the trapdoor is this. To measure regression, Galton drew a picture no one had quite drawn before: children's heights against parents' heights, a cloud of points, and through the cloud a straight line — the line that best summarized the crowd, the line the cloud leans on. And a line like that does something a mere summary has no business doing: *given a new father, it answers*. Tell me his height and I will read off, from the slope, the best guess for his unborn son's — not a certainty, a best guess, the centre of the bell of possibilities. Then came the thunderclap, which arrived, Galton says, while he sheltered from a shower of rain: the line runs both ways. Exceptional fathers have less exceptional sons — and exceptional sons have less exceptional fathers, which no theory of inheritance can explain, since sons do not cause their fathers. The regression, therefore, wasn't biology at all. It was geometry — a property of any two quantities that share a bell and are imperfectly linked, whether or not either causes the other. What Galton had found, and by 1888 had learned to reduce to a single number, was *correlation*: the degree to which

two things vary together, cause be damned. Forearm and height. Parent and child. Barometer and storm. Yesterday's word and today's. It is very hard, from inside a civilization built on this idea, to feel how new it was — a way of extracting prediction from data *without passing through explanation*, a mathematics that says: I do not know why these move together, I decline even to ask, but I have measured how much, and I will bet accordingly. Galton, who had a weakness for the grand phrase and had earned it, wrote of the error law that the Greeks would have deified it had they known of it — the "supreme law of Unreason," he called it, whereby out of a mob of chaotic elements an orderly shape emerges the moment you take them in hand together. And around the new mathematics a new profession promptly crystallized: Karl Pearson — who gave correlation its working formula and the letter r , invented the chi-squared test, co-founded the journal *Biometrika* when the Royal Society got sniffy, and in 1911 established at University College London the world's first university department of statistics, funded by Galton's will. (Pearson's philosophical manifesto, *The Grammar of Science*, was the first book young Einstein's reading club in Bern assigned itself, which is either a coincidence or the whole thesis of this book in one sentence.)

And now the part that must be said plainly, because Galton and Pearson said their part plainly and it would be cowardice to mumble ours. The engine driving all of it — the peas, the threepenny measurements, the correlation table, the endowed chair — was Galton's conviction that if talent and temperament are heritable like height, then humanity could and should breed itself better, encouraging the able to multiply and the "unfit" to refrain. He coined the word for it in 1883: *eugenics* — well-born. The chair Pearson held was the Galton Chair of Eugenics; the laboratory was the Eugenics Record Office before it was anything else; the founders of modern statistics founded it, in large part, *as* the instrument of that project. And the project went where such projects go. In the twentieth century, eugenics acquired laws: sterilization statutes across the United States from

1907, upheld by the Supreme Court in 1927 in *Buck v. Bell*, where Justice Oliver Wendell Holmes disposed of a young Virginia woman's bodily integrity with the sentence "Three generations of imbeciles are enough" — tens of thousands of Americans were sterilized under such laws — and the German statute of 1933 drew openly on the American models, and where that road ended the reader knows without my help. It is no defence, but it is a fact, that the mathematics itself kept warning them: regression to the mean is precisely the announcement that the exceptional do not breed true, that nature herself perpetually dismantles the aristocracies of every trait — Galton spent years wrestling with this, treating as an obstacle to his programme what was actually its refutation. The tools outgrew and outlived the cause they were built for; correlation and regression went on to underwrite medicine, agriculture, economics, and every safety trial that has ever saved your life, and University College has lately taken the founders' names off its lecture theatres, and both facts belong in the record. The axemaker's gift, as ever: the same blade, whichever way you swing it. We keep the mathematics. We had better keep the memory of what it was first swung at, too.



So now let us gather up the thread and follow it home, because everything this chapter has been assembling — the least-squares fit, the bell of errors, the firing table, the regression line — snaps together, in our own day, into a single machine, and you have already met it. It wrote my obituary in the Introduction.

Begin with the innocent question the whole century of this chapter kept asking: what is the best line through a cloud of points? Legendre answered it with a rule; Gauss told us why the rule is right; Galton discovered that the line, once drawn, *predicts*. Now notice something about how Legendre's rule feels from the inside. The sum of squared errors is a number that depends on where you put the line — tilt the line, the

number changes — so imagine a landscape: every possible line a location, the total error at that line its altitude. For the simple straight-line fits of the nineteenth century, that landscape is a single smooth bowl, and finding the best line means finding the bottom of the bowl, which calculus can do in one step, in closed form, by candlelight; that is exactly what made least squares a clerk's tool rather than a philosopher's. But suppose the bowl were not a bowl. Suppose the curve you were fitting had not two adjustable knobs — slope and intercept — but thousands, and the landscape of error had thousands of dimensions and no tidy formula for its floor. What then? Then you do what a sensible walker does on a foggy mountain: you feel the slope under your boots, take a modest step downhill, and repeat. The method even has a birth certificate in our story — Augustin-Louis Cauchy proposed it in 1847, in a short note motivated, and I ask you to appreciate this, by the drudgery of *astronomical orbit calculations*: the direct descendants of Gauss's Ceres problem, grown too fat for closed-form solution. The moderns call it gradient descent — rolling downhill on the error surface — and it is, when the fog lifts, the whole of the secret. There is no second secret.

Because here — shorn of its own jargon, which exists partly to keep the mystery margin healthy, a habit of the knowledge trades you will hear much more about in a later chapter — is what the celebrated machine learning of our era is. Take a mathematical function with a great many adjustable knobs. Show it examples: situation in, outcome out — elevation and splash, father and son, yesterday and today. Define the error: how far the function's guesses fall from what actually happened, totted up over all the examples, misses punished. That error is an altitude; the knobs are your position; now walk downhill, and keep walking, a billion small steps, until the function's guesses stop improving. Legendre would recognize every element — his residuals, his minimization; Gauss would recognize the logic that justifies it; Galton would recognize what comes out, because what comes out is his regression line, promoted: a rule that has passed

through no explanation whatever, that knows nothing of causes, but that has measured, from a mountain of past cases, how things vary together — and will therefore bet, given a new case, on where the next point lands. Fitting a curve to data and learning from experience are not analogous processes. They are, in this trade, the same process, under two names bestowed by two centuries.

And the talking machine? The talking machine is this, and I mean *is*: the largest curve-fit ever attempted by our species. The situations shown to it are stretches of human text — the loom-threads of this book's title, everything we ever wrote down and fed it. The outcome to be guessed is simply the next word. The knobs — the count has lately passed the trillion — are nudged, downhill on the error surface, step after step across months of computation that would have consumed the lifetimes of every human computer who ever sharpened a pencil at Aberdeen, until the machine's surprise at what we actually said next has been ground down as far as it will go. That is the whole of the training, and you should hear in it the oldest cadence in this chapter. A cloud of past observations. A curve drawn through them, offended as little as possible. A reading taken off the curve at a point never observed. *Where will the shell land?* The gunner asked it of the shot, Halley of the heartbeat, Quetelet of the murder rate, Galton of the unborn son; the machine asks it, several hundred times a second, of the sentence you are typing — and answers so well that the answer can pass for conversation, for scholarship, for my own obituary. Why guessing the next word well should shade into something that behaves like understanding is a genuine and beautiful question, and it belongs to a later chapter, where a man at Bell Labs will ask it properly. What belongs here is the pedigree: when the machine completes your sentence, you are watching Piazzzi's lost planet come out of the sun's glare, one word at a time, exactly where the curve said it would.

But carry one more thing forward, because you will need it near the end of this book, and it was Galton who left it on the table. A prediction read

off a fitted curve is not any of the observations. It is their centre — the mean of the bell of possibilities, the most probable continuation, the smoothed middle of everything the cloud contains. That is its power: the wild individual errors cancel, and what remains is steadier than any witness. And that is also its levy: whatever in the data was extreme, singular, unrepeatable — everything, in short, that was somebody in particular rather than everyone on average — is precisely what the curve is built to shave away. Regression toward mediocrity is not a flaw the machine's builders have yet to iron out. It is the operating principle, the very thing Legendre's squares are minimizing toward; the machine speaks from the centre of the distribution because the centre is what it is *for*. What that does to a civilization that starts taking its sentences, its schoolwork, its pictures, and perhaps its opinions from the smoothed middle of its own past — whether we are about to average ourselves the way Quetelet averaged the Scots — is a question I am deliberately leaving loaded on the table, safety-catch off, for the closing chapters. When you meet Galton's ghost there, you will know where the body is buried. It is buried here, in the mathematics of the bell, which gives with one hand a truth no individual possesses and takes with the other the individuality of everyone it summarizes.

There remains only the doorstep. A single fitted curve, however vast, draws straight-ish conclusions from the world; the astonishment of our era arrived when engineers learned to *stack* the fits — to feed the output of one fitted function into another and another, layer upon layer, each layer bending what the last one guessed, the whole tower trained downhill as one — and discovered that a sufficient stack of Legendre's dull little minimizations can learn to see a face, parse a sentence, and complete a thought. Why anyone imagined stacking them in the first place is not a story from the ledgers of astronomers or the proving grounds of artillerymen at all. It begins in a professor's laboratory in Bologna in the 1780s, with thunderstorms, a brass hook, and the severed leg of a frog that

kicked when it had no earthly business kicking — and it is, I'm afraid, the next chapter. The gunners taught the curve to predict. It took the anatomists to teach it to twitch.

CHAPTER FOUR

Nervous Energy

The brain is waking and with it the mind is returning. It is as if the Milky Way entered upon some cosmic dance. Swiftly the head-mass becomes an enchanted loom where millions of flashing shuttles weave a dissolving pattern.

— CHARLES SHERRINGTON, *MAN ON HIS NATURE*
(1940)

One evening in the late 1780s, in a house in Bologna, a professor of anatomy was skinning frogs. Not the opening you expected for a book about artificial intelligence, but bear with me, because almost everything that follows — the thinking machine, the talking machine, the machine that told the world I was dead — begins with a dead frog on a copper hook, and a man who could not explain why its legs kept twitching.

The professor was Luigi Galvani, and the frogs were, as far as anyone could tell, thoroughly deceased. He had a habit of dissecting them for reasons that had nothing to do with the discovery he was about to blunder into — something to do with the effect of electricity on animals, electricity being the great fashionable mystery of the age, the thing every gentleman with a spare room and a static machine was busy sparking at his friends over the port. There are several versions of what happened next, and Galvani himself told the story more than one way, which is your first clue that we are in the presence of a genuine event, because the tidy ones are usually invented. In the version most often repeated, an assistant's scalpel happened to touch an exposed crural nerve at the exact moment a nearby machine threw off a spark, and the dead leg kicked as though the frog had views on the matter. In another version — the one Galvani liked, because it pointed where he wanted to go — he had hung a row of prepared frogs' legs from brass hooks along an iron balcony railing to dry, and noticed that they twitched whenever the leg swung against the iron, spark or no spark, sky clear, no machine in sight. No storm. No visible cause. Just the metal, the flesh, and the kick.

Now, a lesser man would have shrugged and gone to dinner. Galvani was gripped, in the way that changes a life and, occasionally, a civilization. Here was dead tissue moving. Not being moved — *moving*, on its own account, as if some vital spark had been sleeping in the muscle all along and the touch of metal had simply woken it up. He spent the next decade at it, and in 1791 he published his conclusion in a Latin treatise with the sort of title they went in for then — *De viribus electricitatis in motu musculari commentarius*, a commentary on the forces of electricity in muscular motion — and the conclusion was this: living things generate their own electricity. Not the sky's electricity, not the machine's. Their own. He called it *animal electricity*, and he located its reservoir in the nerves and its power in the muscle, and he thought — you can hardly blame him — that he had found the physical basis of life itself. The mysterious fluid the

philosophers had been calling the animal spirits, the thing that made the difference between a corpse and a man, turned out to be nothing more nor less than electricity, bottled in flesh.

He was wrong. He was also, in a way that took a hundred and fifty years to become clear, righter than he could possibly have known. And that gap — between the beautiful wrong idea and the useful right one — is more or less the whole subject of this chapter.



Because there was a man across the water in Pavia, another professor, who read Galvani's treatise and did not believe a word of it. His name was Alessandro Volta, and he was that most dangerous of critics, the kind who admires you. Volta repeated Galvani's experiments — this is what you did in those days, you didn't peer-review a man, you actually tried to make his frogs kick — and the frogs kicked, all right, but Volta drew the opposite lesson. The electricity, he decided, was not coming from the animal at all. It was coming from the *metal*. Look, said Volta, at what actually happens: you get the best twitch when you touch the nerve with one metal and the muscle with another — copper and iron, say, or silver and zinc — and join them. Two different metals and something wet in between. The frog isn't the battery. The frog is the *detector*. It's a very sensitive, very dead little voltmeter, twitching to register a current that the two metals are generating between them, using the frog's own juices as the damp filling in the sandwich. The life had nothing to do with it. You could, in principle, throw the frog away.

And that, my friends, is one of the great quarrels in the history of science, and like most great quarrels it was so productive precisely because both men were partly right and neither would give an inch. Galvani, stubborn and by now getting old, kept hunting for animal electricity without any metal at all — and, to his credit, he found it, in an experiment where a nerve touching a cut muscle produced a twitch with no dissimilar

metals involved whatever, which was, though nobody understood it for a century, the first observation of the actual electrical signal the nervous system runs on. Volta, meanwhile, kept chasing the metals. And here is where the quarrel turns the corner and walks straight into the modern world, because Volta, to prove that you didn't need an animal to make electricity, set out to make electricity without one. If it was the metals doing the work, then a stack of metals should do the work with no frog required. So around 1799 and 1800 he piled up little discs — zinc, then a scrap of cloth soaked in brine, then silver, then zinc again, then brine, then silver, over and over, a tower of metal and salt water — and the tower produced a steady, continuous current, the first that human beings had ever had on tap. He called it a *pile*. We call it a battery. It is, in a real sense, the frog with the frog taken out — the frog abstracted, cartooned, reduced to its electrical essentials and rebuilt in hardware. Remember that phrase, because we are going to do the same thing to the neuron before this chapter is finished, and it is going to work just as well and mean just as little.

The quarrel gave both men a kind of immortality that neither would have chosen. We measure electrical potential in *volts*, after the man who insisted there was no life in it. And when we say a thing has been *galvanized* into action, we are speaking the name of the man who was sure there was. The language, with its usual tactlessness, split the prize down the middle and gave each of them the half they'd have least wanted.



But it was Galvani's half — the wrong, thrilling, living half — that caught the public imagination, because the public has never in its history preferred the correct explanation to the exciting one, and I say that with affection, being the public myself most days. The idea that electricity was the secret fire of life, that a dead thing might be jolted back to motion by the right spark in the right place, was simply too good to leave in the

laboratory. It escaped, as these things do, and it went straight to the most theatrical member of the family.

Galvani had a nephew, Giovanni Aldini, who inherited his uncle's ideas and precisely none of his uncle's reticence. Where Luigi had agonized over frogs, Aldini took the show on the road, and the show got bigger as it travelled, from frogs to oxen to the severed heads of oxen to, eventually and inevitably, a man. On the 18th of January, 1803, a young man named George Forster was hanged at Newgate for the murder of his wife and child by drowning, and — this being the arrangement of the day, the bodies of the hanged belonging by law to the anatomists — his corpse was carried, still cooling, the short distance to where Aldini was waiting with his rods and his battery — a house near the prison, says the *Newgate Calendar*; the theatre of the College of Surgeons, says the College itself; the accounts disagree, and after two centuries I see no reason to pick a winner. What happened next was reported, breathlessly, all over London. Aldini applied the conducting rods to the dead man's face, and the jaw quivered, the muscles of the face writhed into a grimace, and one eye opened. He applied them elsewhere, and the right hand clenched and rose; he applied them to the legs, and the legs kicked. To the crowd of surgeons and gawpers packed into that room it looked, for all the world, as though the murderer were about to sit up off the slab and finish what the rope had started. The official account records, with magnificent understatement, that one onlooker was so alarmed he died shortly afterwards.

Now, nothing was reanimated. Forster was as dead at the end of the afternoon as he had been at the beginning; what Aldini had demonstrated was that dead muscle will contract when you run a current through it, which is exactly what his uncle's frog legs had been saying all along, only louder and in English. But the meaning of a demonstration is never the same as its content, and what the crowd took away was not a fact about muscle physiology. It was a possibility. The line between the living and the

dead, which every religion and every law had treated as absolute, suddenly looked like a switch — and a switch can be thrown the other way.

The culture went to work on the idea the way cultures do, which is to say it made art out of its own anxiety. Thirteen years after Newgate, in the wet summer of 1816 — the year without a summer, as it happened, a volcano in the Dutch East Indies having thrown enough ash into the sky to cancel the season across Europe, which is a Connections story for another day — a party of English exiles was trapped indoors by the rain on the shore of Lake Geneva, and to pass the evenings Lord Byron proposed that each of them write a ghost story. Most of them gave up. One of them, a girl of eighteen named Mary Shelley, could not think of anything, until one night she lay awake and saw, as she later put it, "the pale student of unhallowed arts kneeling beside the thing he had put together." In the introduction she wrote for a later edition, she remembered the talk that had fed the dream — talk of the experiments of Dr Darwin (Erasmus, the grandfather, not Charles), and of the possibility that "perhaps a corpse would be re-animated; galvanism had given token of such things." There is the word, dropped into the birth of the most durable myth of the industrial age. *Galvanism*. Frankenstein's monster is Aldini's murderer with the story allowed to continue past the point where the real one stopped, and the whole modern dread of the made mind that turns on its maker — a dread we are currently living inside, and paying good money to a great many earnest committees to worry about on our behalf — was assembled, in a rented villa, out of a rainy fortnight and a dead frog's leg. The axemaker's gift again: a professor skins an amphibian to understand life, and two centuries downstream we are legislating about killer robots. Nobody plans these things. They arrive sideways, wearing a disguise, and the disguise this time was a monster with bolts in its neck.



Let us leave the theatre and go back to the laboratory, because while the public was thrilling to reanimated corpses, the quiet men with the good instruments were slowly, unglamorously, finding out what a nerve actually *is* and what it actually *does*, and the answer was going to be far stranger than any monster.

The first great surprise concerned speed. Everyone had assumed — it seemed too obvious to question — that a nerve signal travelled instantaneously, or as near to it as made no difference. You touch a hot stove; you snatch your hand back. Surely the message from finger to brain to muscle flashes through you at the speed of the electricity in a wire, which is to say at very nearly the speed of light. The great German physiologist Johannes Müller, the most authoritative voice of the age, taught precisely this, and went further: the speed of nerve conduction, he declared, would very likely never be measured by man, because it was simply too fast, on the order of the speed of light, and the distances inside a frog, let alone a philosopher, were far too short for any instrument to catch the delay. It belonged, in effect, to the realm of the immeasurable — a polite scientific way of saying the realm of the soul.

Müller had a pupil named Hermann von Helmholtz, and Helmholtz did not believe in the immeasurable. In 1850, using a frog's leg, a nerve, and an apparatus of his own devising cunning enough to catch intervals of thousandths of a second, he did the thing his teacher had said could not be done. He stimulated a nerve at one point, then at a point further along, and measured the difference in when the muscle twitched. And the number he got was not the speed of light. It was not anything close to the speed of light. It was about twenty-seven metres a second. Twenty-seven metres a second — slower than a falling stone, slower than a modern sprinter would think acceptable, a speed you could very nearly outrun. The thought in your head, the flinch in your hand, the whole electrical commerce of the self, was ambling along your nerves at the pace of a country bus.

Stop and feel the size of that for a moment, because it is one of the genuinely vertiginous facts in the history of ideas. The mind had always been imagined as instantaneous because it was imagined as immaterial, outside time, a spark of the divine. Helmholtz put a clock on it and found it dawdling. If thought took *measurable time* — if it could be early or late, if it had a *speed* — then thought was a physical process happening in physical stuff, subject to physics, in principle no more mysterious than the boiling of a kettle, only much better organized. He had, without any fuss, moved the mind out of the chapel and into the workshop. And once a thing is in the workshop, someone will eventually try to build a copy of it. That is what workshops are for.



If the mind was in the workshop, the obvious next question was: what is it *made of*? What are the parts? Take the machine apart, as the axemaker taught us to do — that reductionist habit, the one that chips pieces off the universe to see how they fit — and see what falls out.

The trouble was that nobody could see the parts. Nervous tissue under the microscope was a grey, translucent, hopeless tangle, a plate of very fine, very transparent spaghetti with no beginning and no end, and no stain then known would pick out one strand from the mess. Enter, in 1873, an Italian physician named Camillo Golgi, working — the detail is too good to omit — in the kitchen of a hospital where he was the resident medical officer, by candlelight, in his spare time. Golgi discovered, more or less by accident, that if you soaked nervous tissue in silver nitrate after hardening it just so, the silver would, for reasons still not entirely understood, blacken a few cells at random and leave all their neighbours perfectly transparent. It was as if, in a photograph of a packed and foggy stadium, one spectator in a hundred were suddenly rendered in sharp black ink, every hair distinct, while the rest dissolved to nothing. For the first time you could see a single nerve cell whole — its body, its branches, the long trailing fibre —

standing out alone against the emptiness. Golgi called it *la reazione nera*, the black reaction. We call it the Golgi stain, and it is still in use, and it was one of those tools, like the telescope, that does not so much answer a question as reveal that the question exists.

Golgi looked at what his own stain revealed and drew the wrong conclusion, which happens more often in this book than the textbooks like to admit. He saw the branches of neighbouring cells interlacing, mingling, apparently fusing, and he concluded that the nervous system was a single continuous web — a *reticulum*, a net in which the "cells" were merely thickenings, all of them physically joined into one gigantic seamless tissue through which the nervous force flowed like water through a sponge. It was a reasonable reading of a beautiful and confusing picture, and Golgi held to it, fiercely, for the rest of his life.

The right conclusion was drawn, using Golgi's own stain, by a man who could not have been more different from him, and whom Golgi would come to regard with something close to loathing. Santiago Ramón y Cajal was a Spaniard from the Aragonese back country, the son of a poor country surgeon, and by his own cheerful account a nightmare of a boy — a truant, a fighter, once locked up by the village authorities at ten for demolishing a neighbour's gate with a homemade cannon. He wanted to be an artist; his father, appalled, marched him round graveyards and dissecting rooms to make a doctor of him instead, and in one of the great ironies of the discipline it was exactly the thwarted artist's eye — the draughtsman's patience, the hand that could hold a shape in the mind and put it on paper — that made Cajal the greatest observer the brain has ever had. He took up Golgi's stain in the 1880s, improved it, and then simply *looked*, longer and harder and with a better eye than anyone before or since, and he drew what he saw. His drawings hang in art galleries now, and deserve to, but they were never art; they were evidence, and the evidence said something Golgi's net could not accommodate.

The cells, Cajal insisted, were *not* joined. They came achingly, intimately close — the branches of one reaching toward the body of the next like fingers not quite touching in a fresco — but they did not fuse. Each nerve cell was a separate individual, a discrete unit, bounded by its own membrane, a citizen and not a mere thickening in a common jelly. And the signal did not flow through them like water through a sponge; it passed *from one to another*, in one direction, across the tiny gap between, like a message handed over rather than a current conducted. This was the *neuron doctrine*, and it is the foundational fact of everything that follows in this book, so let me put it as plainly as I can: the brain is not a single thing. It is a stupendous number of separate little things, wired together, talking to each other across gaps. It is not a substance. It is a *network*. Hold on to that word too.



The clash came to a head in the most exquisitely awkward setting the scientific world can arrange, which is the Nobel ceremony. In 1906 the Prize in Physiology or Medicine was awarded, jointly, to Golgi and Cajal — to the man who invented the stain and the man who saw what it meant, which sounds like a tidy piece of committee justice until you remember that the two men had spent twenty years publicly calling each other wrong, and still did. They came to Stockholm. They gave their Nobel lectures on consecutive days. And Golgi — old, eminent, and immovable — stood up and used his lecture to defend the reticular theory, the continuous net, the very idea that Cajal's work had by then comprehensively demolished, taking the opportunity of the world's highest scientific honour to insist, in front of everyone, on the thing he had got wrong. Cajal, sitting there, was mortified and furious in roughly equal measure, and said so, in a memoir, later. Two men, one medal, sharing a stage and a prize and absolutely nothing else, each convinced the other was blind. It is my favourite Nobel story, because it is so honest about how knowledge actually advances — not

by consensus, not by the graceful yielding of the old to the new, but by two stubborn men glaring at the same photograph and seeing different worlds, until one of them dies and the students quietly go with the other. Cajal was right. But Golgi's stain is why anyone could tell.

There was one piece still missing, and it was the gap itself. Cajal had inferred the gap between neurons but could not see it — it was far below the resolution of any light microscope, and would not be directly imaged until the electron microscope arrived decades later. It was named before it was seen, which is a very characteristic thing for science to do. The English physiologist Charles Sherrington — the same Sherrington whose enchanted loom sits at the head of this chapter, and no accident that the man who first understood the junction reached, when he wanted to describe the whole, for the image of a weaving machine — needed a word for the point of contact across which one neuron influences the next, and around 1897, on the advice of a classicist friend who suggested a suitable bit of Greek, he coined *synapse*, from the Greek for "to clasp together." The clasp. The gap that is not quite a gap. Sherrington spent his career working out the logic of what happened there, and the logic was the thing that mattered, because it turned out that a neuron does not simply pass its neighbour's message along. It *adds up*. It gathers the nudges arriving across all its synapses — this one urging it to fire, that one urging it to keep quiet — and if the sum crosses some threshold, it fires; and if it doesn't, it doesn't. Excitation and inhibition, pooled and totted up, all or nothing, yes or no.

Now stand back and look at what four generations of frog-slicers and silver-stainers had, between them, assembled, without any one of them meaning to build the thing they were building. A brain is a vast network of separate units. Each unit receives signals from many others across little gaps. Each unit adds up the signals it receives, the yeses and the noes. If the total crosses a threshold, the unit fires and sends its own signal onward; otherwise it stays silent. That is not a description of a mysterious vital fluid.

That is not the breath of God. That is a *wiring diagram*. And a wiring diagram, unlike a soul, is the sort of thing an engineer can look at and say the four most dangerous words in the history of technology: *I could build that*.

★

Two of them said it, in 1943, and one of the two is the reason I am telling this story at all, because if ever a life proved that the future arrives from the least likely direction, it was the short, sad life of Walter Pitts.

Pitts was born in 1923 in Detroit, into a rough family in a rough part of town, the son of a father who wanted him to leave school and get to work and had no patience for a boy who did nothing but read. The origin story — and I flag it as an origin story, because the best-documented lives acquire legends and this one is better documented than most — runs like this. One day, aged twelve, running from a gang of neighbourhood boys, Pitts ducked into the public library to hide. He stayed hidden for three days. And to pass the time he took down, and read, cover to cover, all three volumes of Bertrand Russell and Alfred North Whitehead's *Principia Mathematica* — the towering, forbidding attempt to derive the whole of mathematics from pure logic, a work that has defeated a very large number of professional adults and which a hungry twelve-year-old apparently ate for breakfast in a library he was hiding in. Not only did he read it; he found what he took to be errors, and he wrote to Russell, in England, to point them out. And Russell — this is the part that would be too much if it weren't attested — wrote back, took the boy seriously, and, it's said, invited him to come and study at Cambridge, an invitation a penniless twelve-year-old in Detroit had no earthly way to accept.

Some years later Russell came to lecture at the University of Chicago, and there, sitting in on the lectures, was Pitts — by now a teenage runaway, homeless, drifting around the university with no enrolment and no money, doing menial work when he could get it, and knowing more logic than most of the faculty. He fell in with the sort of people who collect

strays of genius, and eventually, in 1942, he was introduced to a man named Warren McCulloch. McCulloch was everything Pitts was not: forty-three, established, a neurophysiologist and psychiatrist, a big bearded poetry-writing eccentric with a house and a family and a settled place in the world. He was also chewing on exactly the problem Pitts was equipped to solve. McCulloch had been brooding for years on whether the all-or-nothing firing of Sherrington's neurons — the yes-or-no, the threshold, the sum of excitation and inhibition — might be the same thing, at bottom, as the true-or-false of formal logic. What Pitts brought was the *Principia*: the machinery for treating true-and-false, and-and-or-and-not, as a rigorous calculus. McCulloch, learning that this astonishing homeless teenager had nowhere to live, took him into his own home, and there, often late at night, the middle-aged professor and the vagrant prodigy worked it out together.

What they published in 1943, in a mathematical biology journal almost nobody read at the time, was called "A Logical Calculus of the Ideas Immanent in Nervous Activity," and it is one of the founding documents of the age we are living in, though it contains not a single line of anything we would call a computer program, because there were, as yet, essentially no computers. What McCulloch and Pitts did was to take Cajal's neuron and Sherrington's threshold and do to them exactly what Volta had done to Galvani's frog: throw away everything living and keep the logic. Their "neuron" was a cartoon — a simple unit with inputs, each input either exciting or inhibiting, a threshold, and an output that was on or off, one or zero. No chemistry. No ion channels. No mess. And then they proved a theorem that ought to have been on the front page of every newspaper in the world, only nobody yet understood what it meant: a network of these cartoon neurons, wired up correctly, could compute *any* proposition that logic could express. Any statement you could frame in the calculus of true and false, and, or, not, if-then — any of it — could be carried out by some arrangement of these idealized on-off cells. The brain, at least the logical

skeleton of it, was a computing machine, and — the deeper and stranger claim — a computing machine, if it were built out of the right kind of logic gates, might in that precise sense be a brain. The wiring diagram had become a proof.

I want to be careful here, because it is easy to make this sound cleaner than it was. McCulloch and Pitts had not built anything, and their cartoon neuron leaves out very nearly everything a real neuron does. But they had done the thing that matters most in the whole of this book, more than any invention: they had made the *identification*. They had said that thought and computation are, at some fundamental level, the same activity — that what the neuron does when it adds up its inputs and decides whether to fire is not *like* logic, it *is* logic — and once that identification is made, once the mind is admitted to be a kind of calculation, the road runs straight, though long and through two terrible winters, to a machine on your desk that argues with you about poetry.

As for Pitts, the story does not end well, and I would be cheating you if I pretended otherwise. He went to MIT, fell in with the great cybernetician Norbert Wiener, was for a while the golden boy of the whole new science of thinking machines. Then Wiener, for reasons tangled up in his own family quarrels, cut Pitts off completely and without explanation, and Pitts — who had rebuilt his whole life around that mentorship — never recovered. He drank. He withdrew. He is said to have burned his unpublished work, including a doctoral thesis, in one of those acts of self-erasure that break the hearts of everyone who knew the man. He died in 1969, of the consequences of drink, at the age of forty-six, holding no degree of any kind, having read the *Principia* in a library he was hiding in and helped, from that beginning, to invent the intellectual foundation of artificial intelligence. Keep a thought for him: genius and misery keep company far more often than the inspirational posters admit.

The cartoon neuron of 1943 could, in principle, compute anything logic could state — but only if you wired it up by hand, in advance, to do that one thing. It could not *learn*. And learning, the changing of a mind by its own experience, is the whole point; a network that has to be told, synapse by synapse, exactly what to be is just an unusually complicated way of writing down what you already knew. Where was the learning going to come from?

The answer had, characteristically, already been half-guessed by looking at the biology. In 1949 a Canadian psychologist named Donald Hebb published a book, *The Organization of Behavior*, containing a single idea so simple and so fertile that it now underlies essentially every learning machine on Earth. Hebb proposed that when one neuron repeatedly takes part in making another neuron fire — when cell A is, again and again, one of the voices that successfully pushes cell B over its threshold — then the connection between them *strengthens*, so that A's vote comes to count for more in B's future decisions. Do a thing together often enough and the wire between you thickens. The popular slogan for this — "cells that fire together, wire together" — was actually coined by somebody else, decades later, and Hebb never put it so snappily, but the compression is fair and I'll use it with the wink it deserves. The importance of Hebb's rule is that it locates *learning in the connections*. Not in the cells, which are dumb little adding machines all much alike, but in the adjustable strengths of the wires between them — in what an engineer would call the *weights*. Experience does not rewrite the neurons. It re-weights the network. Remember that, when we get to the end of this chapter and I try to tell you where the magic lives: Hebb told us in 1949, and it was in the wiring all along.

Now we have all the pieces. A unit that adds up its inputs and fires past a threshold — McCulloch and Pitts. Connections of adjustable strength — Hebb. And, inherited from the previous chapter, from Gauss and Legendre and Galton and the whole apparatus of learning-from-error that they built to predict where the artillery shell would land, a way of *tuning* those

strengths by nudging them, over and over, a little in the direction that reduces your mistakes — what that chapter called, and what everyone now calls, gradient descent, the slow rolling downhill toward the least wrong answer. Put the three together — the threshold unit, the adjustable weights, the downhill tuning — and you have the recipe for a machine that improves itself by practising. All it wanted was somebody with enough nerve to build it and enough showmanship to make the world care. It got both, in the same man, and the man was trouble.



Frank Rosenblatt was a psychologist by training, not an engineer, which turned out to matter, because he was interested in the brain as a thing that *learns to perceive* rather than as a thing that does logic, and perception — recognizing a shape, telling an A from a B, seeing a face in a crowd — is exactly the sort of squishy, statistical, no-clean-rules task that formal logic is worst at and that living brains do without breaking a sweat. In 1957 and 1958, at the Cornell Aeronautical Laboratory, funded by the US Office of Naval Research, Rosenblatt built a machine he called the Perceptron.

It was, in its first incarnation, less a program than a room full of apparatus — an array of four hundred photocells, a bank of potentiometers whose settings were the adjustable weights, motors to turn the knobs — and its trick, its one genuinely marvellous trick, was that it learned. You showed it cards; it guessed; you told it whether the guess was right; and it reached in and adjusted its own weights, by exactly the downhill nudging we've just described, so that next time it would be a little less wrong. Nobody programmed it to recognize a shape. It *worked the shape out for itself*, from examples — the artillery problem of the last chapter, solved not with a pencil but with photocells and motors, learning to predict not where a shell would land but which letter it was looking at. It was, I do not think this is too strong, the first machine that visibly, undeniably, taught

itself something its makers had not put into it. For anyone who had read McCulloch and Pitts and Hebb, here was the theory made of metal and clicking to itself in a laboratory in upstate New York.

And Rosenblatt, God help him, could *talk*. He briefed the press, and the press did what the press does, which is to take the most electrifying sentence in the room and print it without the surrounding caution. In July 1958 the New York Times carried a report — I have read it, and it is a museum piece — announcing that the Navy had revealed “the embryo of an electronic computer” which it expected “will be able to walk, talk, see, write, reproduce itself and be conscious of its existence.” From a machine that could, on a good day, tell a printed A from a printed B, and not always that. It is the exact template of every AI announcement you have read in your own newspaper this year: a genuine, remarkable, narrow result, wrapped in a prophecy of imminent godhood, delivered by someone who half believes it and wholly needs the funding. We have been doing this since 1958. We will be doing it long after the machine has, in fact, learned to write, at which point we will move the goalposts and complain that it still can’t love. That, too, is the pattern.



Overpromising has a price, and the bill for Rosenblatt’s came due in 1969, delivered by a man who had sat two rows over from him in high school.

Marvin Minsky had been at the Bronx High School of Science with Rosenblatt; they knew each other; and Minsky had by now become, at MIT, the most powerful figure in the rival camp — the camp that believed the road to artificial intelligence ran through logic, symbols, and explicit reasoning, the grand tradition of McCulloch’s true-and-false rather than Rosenblatt’s messy statistical learning. The two approaches were competing for the same limited pool of money, and Minsky, with his colleague Seymour Papert, wrote a book, called simply *Perceptrons*, that set out to es-

tablish, with full mathematical rigour, exactly what Rosenblatt's machines could and could not do. And what it proved was devastating in its narrowness. A perceptron of the simplest kind — a single layer of those adjustable weights — was shown to be incapable, even in principle, of learning certain childishly simple distinctions. The famous one is a function called *exclusive or*, XOR: true if one input or the other is on, but not if both are, and not if neither is. A toddler grasps it; a single-layer perceptron provably cannot, because the yeses and noes it needs to separate cannot be divided by the single straight line such a machine is limited to drawing. It was a real limitation and the mathematics was correct.

It was also, we can now see, a limitation of the *simplest* version only — of a perceptron with no middle layer between input and output. Stack the neurons in layers, put a hidden layer in the middle, and the XOR problem dissolves; a multi-layer network has no such difficulty at all. Minsky and Papert knew this, and said so in the book, but they were pessimistic in print about whether anyone would ever find a way to *train* such deeper networks — the downhill-nudging trick that worked so neatly on a single layer had no obvious way to reach back and adjust the weights buried in the middle, because you couldn't easily say how much each hidden neuron was to blame for the final mistake. The problem was real. The pessimism was the thing that did the damage. The book was read — and, more consequentially, was *cited by people who had not read it* — as a proof that neural networks were a dead end, a toy, a discredited enthusiasm. The funding agencies, already nervous after a decade of walk-talk-see-write prophecies that had not come true, drew the obvious administrative conclusion, and the money for the whole approach simply stopped. This is the event the field now calls, with a shiver, the first AI winter, and it lasted the better part of two decades, and a great many careers froze solid in it.

Rosenblatt did not live to see the thaw. In July 1971, sailing on Chesapeake Bay — on his forty-third birthday, as it happens, the day closing the exact span it had opened — he drowned in a boating accident.

He was forty-three years old, his life's work under a professional cloud, the machine he had built dismissed as a curiosity. He never knew that he had been essentially right, that the winter was temporary, that the thing he saw when he looked at those clicking potentiometers and imagined a machine that would one day see and write was not a delusion but a preview, off by half a century in its timing, which as prophecies go is respectable. The reproach he did not live to answer, and the vindication he did not live to enjoy, both belonged to the same idea, and the idea was only sleeping.



Here is the thing about winters, though, the thing that makes them survivable: somewhere, some small stubborn number of people simply refuse to accept that the sun has gone out. Through the 1970s, when there was no money and no glory and considerable professional risk in saying you worked on neural networks — you might as well have said you worked on alchemy — a scattered handful of believers kept the fire going, mostly by pretending, for grant-application purposes, to be doing something respectable.

In Britain and then Canada there was Geoffrey Hinton, a psychologist's cast of mind in a computer scientist's job, descended, as it happens, from the George Boole whose algebra of true and false runs underneath the whole logical tradition — a nice knot in the thread, the great-great-grandson of Boolean logic spending his life on the anti-logical, statistical, connectionist heresy. In France and then America there was Yann LeCun, who would take the perceptron idea and shape it, cleverly, for images — a network with its wiring arranged to notice small features anywhere in a picture, which he trained to read the handwritten digits on postal envelopes and bank cheques, one of the first neural networks to earn its keep in the actual world. In Switzerland there was Jürgen Schmidhuber, whose lab worked, among much else, on getting networks to remember

things over time. They were not a movement but a diaspora, corresponding, quarrelling, keeping each other's spirits up, and waiting.

And in 1986 they got their answer to the problem that had frozen everything — the problem of how to train the deeper, multi-layer networks that could, in principle, do everything the single layer couldn't. The answer was a way of assigning blame backwards. When the network gets the final answer wrong, you can, it turns out, work mathematically backwards through the layers, computing for each connection — including the hidden ones buried in the middle — exactly how much it contributed to the error, and nudge each one accordingly. The method is called *backpropagation*, back-propagation of errors, and it is, at bottom, the chain rule of ordinary calculus applied with great patience through a many-layered function; pieces of it had been discovered and forgotten more than once before, which is the usual fate of ideas that arrive before their moment. But the paper by Rumelhart, Hinton, and Williams, in *Nature*, in 1986, is the one that made the field sit up, because it showed the thing working: hidden layers learning, by themselves, to represent useful features nobody had specified, the very thing Minsky and Papert had doubted anyone would manage. The theoretical objection that had justified the first winter was answered. Backprop is still, nearly forty years on, essentially how every one of these machines learns. The whole towering edifice of modern AI runs, underneath, on assigning blame backwards and rolling a little way downhill, over and over, a trillion times.

You would think that would have ended the winter for good. It didn't — not yet. Because having the right method is not the same as having the means to use it, and through the 1990s the networks that backprop could train remained, in practice, small and slow and easily beaten at most real tasks by other, less romantic techniques. The theory was sound; the machines were feeble; and so a second winter settled in, milder than the first but longer, through the late 1980s and the 1990s and into the 2000s, during which "neural network" remained a slightly embarrassing phrase in

polite computing society and the faithful remained faithful mostly by main force of character. They kept working. They were, by now, very good at keeping working. It is the least glamorous virtue in science and very nearly the most important one.



What the faithful were missing was not an idea. Every idea was, by the early 2000s, already in place — the threshold unit from 1943, the learning-in-the-weights from 1949, the self-training machine from 1958, the blame-assigned-backwards from 1986. What they were missing was *bigness*. The networks were too small and the examples too few and the arithmetic too slow, and the suspicion had begun to grow, in a few heads, that this was the *whole* problem — that the theory had been right all along and had simply been starved, that if you could only build a network enormously bigger, and show it enormously more examples, and grind through enormously more of the downhill arithmetic, the thing would suddenly, disproportionately, *work*. It was not a fashionable thing to believe. It sounded lazy, almost anti-intellectual: no new clever idea, just *more*. But it was, in the event, correct, and it has a name now — the scale hypothesis — and it is the single most important and least dignified idea in the recent history of the mind. Sometimes the breakthrough is not an insight. Sometimes it is a big enough pile of the thing you already had.

The moment it became undeniable can be dated. In 2012, a neural network built by Hinton and two of his students, trained to sort photographs into a thousand categories in an annual competition that was the field's proving ground, did not merely win. It won by a margin so large that it stopped the argument dead — halving, near enough, the error rate of the best of the traditional methods, in a domain, machine vision, where those methods had been inching forward by fractions of a percent for years. The people in that field describe the day the results came out as a sudden, vertiginous sense that the ground had shifted and was not shifting

back. Within a very short time, everyone who had spent a career on the respectable alternatives quietly retooled, and the whole discipline swung, en masse, back to the heresy it had left out in the cold for forty years. The winter was over. It has not returned.

I am going to leave you in genuine suspense about *how* that 2012 network was trained fast enough to be possible at all, because it is the best twist in the whole book and it belongs to the next chapter, not this one — I will say only that the decisive machine was borrowed from an utterly unlikely source, one that had spent thirty years and untold billions perfecting itself for a purpose that had nothing whatever to do with thought, and of which, as I say, more in the next chapter. For now, the moral is enough: the scattered faithful had been right, the winter had been a failure of hardware and nerve rather than of theory, and the cartoon neuron of Walter Pitts, scaled up beyond anything he or Rosenblatt could have dreamed, had finally been given a big enough loom to weave on.

★

So let me close by keeping the promise I made at the top, and telling you where the magic lives, because the answer is the most deflating and the most wonderful thing in this chapter, and it has been sitting in plain sight since Hebb.

Consider the artificial neuron — the thing, multiplied by the billion, that does the seeing and the writing and the confident lying. It is McCulloch and Pitts' cartoon, barely modified in eighty years. It takes a set of numbers in. It multiplies each by a weight. It adds them all up. If the sum is big enough, it passes something on; if not, it doesn't. That is the entire animal. It is, set beside a single one of Cajal's real neurons — with its thousands of synapses, its living chemistry, its ion channels and its neurotransmitters and its structure so intricate that a century of the finest microscopists have not exhausted it — a childish thing, a doodle, a stick figure of a cell. Every neuroscientist alive will tell you, correctly and with

some feeling, that the artificial neuron is a travesty of the real one, that it throws away almost everything that makes a living neuron what it is. They are right. It is a cartoon of a cartoon: Volta's frog-with-the-frog-removed, done again, to the brain.

And here is the deflating, wonderful punchline: *the cartoon was enough*. You do not need the chemistry. You do not need the living cell in anything like its full glory. You take this laughably simplified stick-figure neuron, and you wire up a very great many of them, and you tune the strengths of the connections between them — the weights, Hebb's weights, the thicknesses of the wires — by rolling downhill against your errors a few trillion times, and the thing learns to see, and to translate, and to write, and to converse, and to do very nearly everything we had spent three thousand years insisting required a soul. Which means the intelligence was never *in* the neurons at all. The neurons are almost nothing; you can replace them with a cartoon and lose almost nothing that matters. The intelligence is in the *connections* — in the pattern of the wiring and the strengths of the links, in the arrangement rather than the parts. It was in the connections when Cajal saw that the cells did not touch but reached; when Sherrington named the clasp between them; when Hebb said that learning lives in the changing strength of the wire; and it is in the connections now, in a machine that has never possessed a single living cell and gets along, apparently, perfectly well without one.

I find that news, on balance, rather marvellous, and I have spent a career being accused of finding everything connected, so you may discount me accordingly. But think what it says. The thing that makes a mind is not some special stuff, not a vital fluid, not animal electricity, not a spark that only the divine can strike. It is a *pattern of relationships* — and a pattern of relationships is precisely the sort of thing that can be lifted off one material and woven onto another, off the wetware and onto the sand, the way a tune can be lifted off a violin and played on a bell. Galvani thought he had found the secret of life in the flesh of a frog, and the secret turned out not

to be in the flesh at all; it was in the wiring, and the flesh was just the first place it happened to be wired. We are, all of us, patterns that learned to hold themselves together — enchanted looms weaving a dissolving pattern, never an abiding one — and you can build the loom out of almost anything, provided you get the threads to cross in the right places.

The frog kicked because of the connection between two metals. The mind wakes because of the connections between ten billion cells. And the machine that told the world I was dead does what it does because of the connections between a great many cartoons of those cells, tuned by rolling downhill, on a loom borrowed — as I keep promising, and the next chapter will finally deliver — from the least likely industry on Earth. Turn the page.

CHAPTER FIVE

Fun and Games

Play is older than culture, for culture, however inadequately defined, always presupposes human society, and animals have not waited for man to teach them their playing.

— JOHAN HUIZINGA, *HOMO LUDENS*

In the spring of 1770, in the palace of Schönbrunn outside Vienna, a Hungarian civil servant named Wolfgang von Kempelen wheeled a large wooden cabinet into the presence of the Empress Maria Theresa and prepared to lie to her. The lie was a beauty. On top of the cabinet sat a life-sized figure in Turkish robes and a turban, holding a long pipe; on the cabinet's surface, a chessboard. Kempelen opened the doors one by one, held a candle behind them so that everyone could see clear through the machinery — brass gears, cylinders, clockwork of the most reassuring kind — closed everything up, wound the mechanism with a great theatrical key, and invited a courtier to play. The Turk played with its left hand (a slip in

the design, most historians think, which Kempelen never bothered to correct — even the best liars have their off days), it rolled its eyes, it nodded when it gave check, and it won. It nearly always won. It would go on winning, off and on, for eighty-four years, which is a remarkably long run for a fraud, when you consider that the average modern investment scam struggles to last five.

Kempelen had built the thing to settle a score with a conjurer. The year before, a Frenchman named Pelletier had entertained the court with tricks of magnetism, and Kempelen — an engineer to his fingertips, the sort of man who designed waterworks and pontoon bridges for a living — had let slip to the Empress that he could contrive something better. Careless talk. The Empress held him to it, and the result was the most successful piece of misdirection of the entire eighteenth century, because of course there was no chess-playing machinery inside the cabinet at all. There was a chess player. A full-grown, flesh-and-blood, usually impecunious and frequently alcoholic chess master, sitting on a sliding seat, shifting himself neatly out of the way as each door was opened in sequence, following the game from below by watching small magnetic pointers that dipped and swung under the board as the pieces (loaded with magnets) moved above, and directing the Turk's arm through a pantograph. He worked by candlelight, in a box, in his shirtsleeves. The clockwork the audience heard was scenery. The whole thing was a puppet, and the puppeteer was the one thing nobody in any audience for eight decades could quite bring themselves to believe in: an ordinary hidden human being. People proposed dwarfs, trained monkeys, evil spirits and — this was the enlightened guess — magnets. Almost nobody said: there's a man in the box. Remember that, because this book will have cause to come back to it. When a machine performs intelligence, the audience will believe almost anything before it believes the simple explanation.

Kempelen himself was faintly embarrassed by the Turk. He called it a *bagatelle*, a trifle, and kept trying to retire it so that he could get on with

what he considered his serious work — which was, and I promise I am not making this up, a genuine talking machine: a contraption of bellows, reeds and a rubber mouth that could actually pronounce words, and about which he published a fat scholarly volume in 1791. Consider the irony, because it is the founding irony of this entire book. The fake thinking machine made him famous across Europe. The real talking machine, the honest one, the one that took him twenty painstaking years, was received as a curiosity and largely forgotten — although not entirely, because decades later the physicist Charles Wheatstone built a reconstruction of it, and a young visitor named Alexander Graham Bell saw that reconstruction, went home determined to make speech out of machinery, and — but that's a thread for another chapter, and it's rude to unravel two at once.

The Turk toured. In Paris in 1783 it played the strongest player of the age, Philidor, who won, though his son claimed long afterwards that the old man called it the most fatiguing game he ever played — sons polish their fathers' anecdotes, so handle that one with tongs. It beat Benjamin Franklin, then the American minister to France, who by all accounts enjoyed himself enormously. And after Kempelen died in 1804, the whole apparatus was bought by a Bavarian showman-engineer named Johann Nepomuk Maelzel — patentee of the metronome, builder of a mechanical orchestra for which Beethoven wrote a battle symphony before the two of them quarrelled and went to law — who understood, as his modern descendants in Silicon Valley understand, that the technology is only half the product. The other half is the show. Under Maelzel the Turk became an industry: advance publicity, ticketed exhibitions, celebrity challengers, and a payroll of hidden operators that reads like a chess directory of the age — Allgaier, Boncourt, Mouret, Lewis, Schlumberger — each of them sworn to secrecy, several of them drunks, one of whom, Mouret, eventually sold the whole secret to a French magazine in 1834 for drinking money. Nobody important believed the magazine. The show went on.

And in 1809, at Schönbrunn again — the palace now full of French officers, the Empire now Napoleon's — the Turk played the most famous game of its career. Napoleon Bonaparte, fresh from Wagram, sat down opposite the automaton and, the story goes, tested it the way he tested everything: by cheating. He made an illegal move. The Turk replaced the piece. He made another; the Turk removed the offending piece from the board and played on. On the third offence the Turk swept the entire board clean with its wooden arm, and the Emperor of the French, delighted, sat back and laughed, then played a proper game and lost it in nineteen moves to the hidden Allgaier. Now, I am obliged to tell you that the earliest full accounts of this scene were written down decades later, and that the details grow with each retelling in exact proportion as the eyewitnesses die off. It is, in short, too good to be entirely true. But it was widely believed at the time and ever since, and for our purposes the belief is the point: by 1809 all Europe was prepared to accept that a machine could not only play chess but could enforce the rules of chess — could stand on principle — against the most powerful man in the world. The idea of the thinking machine was now loose in the culture, eighty years before there was any machinery that could remotely support it. Ideas do that. They arrive before their hardware, travelling light, and wait.

The Turk crossed the Atlantic in 1826 and toured America, where it met two fates, one immediate and one slow. The immediate one: in Baltimore, two boys perched on a rooftop watched the exhibition hall after a performance and saw a man climb out of the cabinet — Schlumberger, the resident operator, coming up for air — and the local press printed the story, and it made no difference whatever, because the public did not want the Turk explained; it wanted the Turk marvellous. The slow fate was a magazine essay. In 1836 a young Richmond editor with a taste for puzzles published "Maelzel's Chess-Player," a long, remorseless piece of deductive reasoning arguing from seventeen observations — the machine's irregular timing, its jerky arm, the curious choreography of the opened doors —

that there must be a man inside. The editor was Edgar Allan Poe. He borrowed some of his arguments (without excessive acknowledgement) from the Scottish scientist David Brewster, and one of his central deductions was magnificently wrong: a true machine, Poe reasoned, being incapable of error, would win *every* game, and since the Turk occasionally lost, it must be human. Hold that thought for a hundred and sixty years; we will be standing in a Manhattan skyscraper watching a world champion who has just lost to a computer insist, on precisely inverted Poe logic, that because the machine played so *humanly* there must be a man in the box. And note what Poe did next: within five years he had invented the detective story — Monsieur Dupin, the locked room, the chain of inference laid out for the reader's pleasure. The method he rehearsed on the automaton became a genre, the genre became Sherlock Holmes, and the fake thinking machine thereby taught us how to write about thinking. Sideways, as always.

The Turk itself ended badly, as stars do. Maelzel died at sea in 1838, broke, sailing home from a disastrous Cuban tour that had already killed Schlumberger by yellow fever. The automaton was auctioned, restored by a Philadelphia physician named John Kearsley Mitchell — Poe's own doctor, small world — and parked in a Philadelphia museum, where on the night of the fifth of July 1854 a fire next door reached the building and burned the Turk to ash. Mitchell's son, the neurologist Silas Weir Mitchell, wrote its obituary, and claimed — with a perfectly straight face, and I pass it on with mine — that through the flames could be heard the creature's last words: "Échec! Échec!" Check. Check. No returns, no explanations, and, in this case, no survivors.



But the Turk had already done its real work, and its real work had nothing to do with entertainment. In the spring of 1819, during the London leg of Maelzel's tour, the machine played two games against a

young English mathematician, and beat him twice. The mathematician was Charles Babbage — you met him at his dinner parties in Chapter Two, demonstrating the Difference Engine to duchesses — and Babbage was not fooled for a moment. His notes make it clear he assumed a hidden player. But being Babbage, he could not leave it there, because the fraud posed an honest question: the Turk was fake, fine — but *need* it have been? Was there anything in the game of chess, in principle, that machinery could not do? He worried at the problem for years and concluded, in his memoirs, that every game of skill could in principle be played by an automaton: that the whole business of positions and moves could be reduced to a mechanical procedure of lookahead and selection. He even sketched the economics, and here the story turns almost painfully modern. Babbage needed money for the Analytical Engine — the real thinking machine, the one with Jacquard's punched cards at its heart — and he seriously considered building a machine to play noughts and crosses, touring it as a spectacle, and using the gate receipts to fund the Engine. He abandoned the idea after concluding that the exhibition business was already saturated (he had made enquiries; the leading attraction of the day was a talking doll, and the margins were poor). Let the record show, then, that the first person to propose funding general-purpose computing research with games revenue was Charles Babbage, in or around 1840. The games industry would eventually take up the suggestion — a hundred and fifty years later, by accident — but I am getting ahead of myself.

What Babbage speculated, a Spaniard built. In 1912 — three years before the word "robot" existed, six before the end of the First World War — an engineer named Leonardo Torres Quevedo constructed a machine that genuinely played chess. No cabinet, no candle, no sweating Frenchman on a sliding seat: an electromechanical device that played the endgame of king and rook against a human opponent's lone king, and mated him, every time, from any position, unassisted. It sensed the pieces electrically, it moved its own rook and king, and if the human tried a Napoleonic illegal

move a light came on; three offences and the machine — I find this detail almost unbearably dignified — simply stopped playing. Torres exhibited it in Paris in 1914 to polite astonishment. He had already invented the Telekino, a radio remote control with which he steered an unmanned boat around Bilbao harbour in 1905, and he would go on to build the aerial cable car that still carries tourists over the Niagara whirlpool, in continuous service since 1916 — possibly the only machine in this book you can actually ride. In 1914 he published an essay on what he called Automatics, arguing that machines could be built with senses, limbs and stored rules, and that they would do many of the things "commonly classed as thought" — the quotation marks are doing honest work there; the paraphrase is close but the Spanish is better. Hardly anyone read it. The chess machine was a footnote in the newspapers, filed under amusements. And in 1951, at a cybernetics congress in Paris, Torres's son Gonzalo wound up his late father's contraption one more time and set it against the grandmaster Savielly Tartakower, who duly, in an exhibition, with a smile on his face, became the first grandmaster in history to lose to a machine. Almost nobody noticed that either. The machine had been ready since 1912. The century had not.



The century caught up in a philosophy journal. In October 1950, the quarterly *Mind* — a periodical not previously known for changing the world — published a paper by Alan Turing called "Computing Machinery and Intelligence," which opens with the question "Can machines think?" and then, in the second sentence, throws the question away as meaningless. Wonderful move. Instead of defining thinking — a bog in which philosophy had been thrashing for two millennia — Turing proposed a game. Put a machine at the end of a teleprinter wire and let it hold a conversation with you. If, after a good long exchange, you cannot tell whether you have been talking to a machine or to a human being — if the

thing on the other end of the wire could pass, let us say, for your aunt: chatty, opinionated, good on family history, occasionally shaky on dates — then the question "but is it *really* thinking?" has no more content than asking whether a submarine *really* swims. (I should confess, since this book has now used the aunt twice and intends to go on using her, that Turing's own version was stranger: it began with a man pretending by teleprinter to be a woman, the machine then substituting for the man — an arrangement scholars have argued over ever since, and I intend to stay well clear of the argument. The aunt is mine. Turing may keep the rest.) The paper made one prediction — that by the year 2000 machines would play the game well enough to fool an average questioner about a third of the time after five minutes, and that "the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted." He was late by about two decades, which by the standards of Chapter Eleven's forecasting hall of shame makes him practically clairvoyant.

But notice — and this is the point of putting Turing in this chapter rather than any other — that when the deepest thinker of the computer age wanted to define intelligence itself, he reached for a *game*. A parlour game, literally: the imitation game was a Victorian party amusement before Turing got hold of it. And when he wanted to demonstrate machine intelligence in practice, he reached for a game again. In 1948, with his friend the economist David Champernowne, Turing wrote a complete chess-playing program called Turochamp — evaluation rules, lookahead, the works — for a computer that did not yet exist. There being no machine to run it, Turing ran it himself: he sat with the rulebook and a pencil and executed his own program by hand, one instruction at a time, about half an hour per move, playing the part of the hardware with a patience his friends found comic and I find rather moving. In 1952 Turochamp-as-performed-by-Turing played a colleague, Alick Glennie, and lost in twenty-nine moves; it did, for the record, beat Champer-

nowne's wife, a beginner, which makes Mrs Champernowne the first human being defeated by a chess program, seven years before there was a computer anywhere that could run one. Savour the symmetry, because you will not meet a better one in this book: the Turk was a machine with a man hidden inside it. Turochamp was a man with a machine hidden inside him. Between those two images — 1770 and 1948, the fraud and its exact inversion — runs the whole story of artificial intelligence.

That same year, 1950, an American engineer at Bell Laboratories published a paper coolly setting out how a computer should be programmed to play chess: how to score a position, how to search the tree of possible moves, why searching everything was hopeless and selective search essential. His name was Claude Shannon, he had already done something far more important that we shall come to, and his chess paper became the founding document of the whole field — every chess machine for the next half century was, in essence, an elaboration of its two strategies. I mention him here only in passing, and with real regret, the way one mentions in passing that the quiet man at the end of the bar happens to have invented the twentieth century. He built toys, too — a mechanical mouse that learned mazes, a calculator that worked in Roman numerals, machines of majestic uselessness — and rode a unicycle down the corridors of Bell Labs, sometimes while juggling. The next chapter belongs to him. He can wait; he was a patient man.



Now the games begin in earnest, and we may as well start with the game everyone condescends to. Draughts — the Americans, who will not be told, call it checkers — was chosen as a target in the early 1950s by an IBM engineer named Arthur Samuel precisely *because* it was humble: simple rules, small board, no aristocratic mystique, and just enough depth to be interesting. Samuel's program, coaxed onto IBM's earliest commercial machines, did something in the 1950s that ought to have made more noise

than it did: it learned. It improved its own evaluation of positions by experience, and — the trick that should have every reader of this book sitting up — it played thousands of games *against itself*, tuning its judgment with each one, no human opponent required. By the late fifties it played considerably better draughts than Arthur Samuel, a fact Samuel reported with the cheerfulness of a man who was, by his own admission, a mediocre draughts player and a superb engineer. When the program was demonstrated on American television in 1956, IBM's share price jumped — the story goes overnight, and the story goes fifteen points, and I note that IBM has never been in any indecent hurry to correct it. In 1959 Samuel wrote up his methods in a paper whose title quietly coined a phrase: "Some Studies in Machine Learning Using the Game of Checkers." Machine learning. You have heard the phrase since. It was born at a draughts board, and the self-play trick born with it will reappear at the climax of this chapter wearing considerably more expensive clothes. (In 1962 IBM's publicity department arranged a match against one Robert Nealey, billed as a champion, whom the program beat; the pedants have since established that Mr Nealey's title was, let us say, generously described. The overclaimed demo was born at the same board. It has also flourished.)

Since we are at the draughts board, let me finish its story, out of order, because it deserves it. The greatest draughts player who ever lived was a Florida mathematics teacher named Marion Tinsley, who between 1950 and 1990 lost a total of somewhere under ten serious games — not ten matches, ten *games* — and who regarded his talent, sincerely and serenely, as a gift from God. In the early 1990s a Canadian team led by Jonathan Schaeffer built a program called Chinook that finally gave Tinsley an opponent worthy of him: the elderly believer versus the machine, genuine respect on — I was about to say on both sides, and who knows, perhaps I mean it. Tinsley won their first match. In the 1994 rematch, after six straight draws, he withdrew feeling unwell; the illness was pancreatic cancer, and he died within the year, unbeaten, the only world champion in

any game to hold off the machines to the end of his natural life. Schaeffer's team then spent thirteen further years and hundreds of computers not merely beating draughts but *solving* it: in 2007 they published the proof that draughts, played perfectly by both sides, is a draw. Think about what that sentence does. A game played by human beings for centuries, a game that consumed whole lifetimes, Tinsley's included, was shown to be — in the mathematician's sense — empty: two perfect players produce nothing, forever. And yet nobody who has understood that result has ever concluded that Tinsley wasted his life. Keep that thought too. The value of play was never in the outcome. This chapter is circling that fact the way the vulture circles lunch.

Chess took longer, because chess is bigger, and the story of its fall is really three stories stapled together: a story about prophecy, a story about hardware, and a story about a man cracking under pressure in a Manhattan office tower. The prophecy first. In 1957 Herbert Simon — a future Nobel laureate, no crank — predicted that a computer would be world chess champion within ten years. Ten years later, in 1967, the strongest program in existence was beaten regularly by decent club players, and the philosopher Hubert Dreyfus had published a gleeful broadside declaring machine chess a dead end, computers unable to beat even a ten-year-old. Whereupon — you could set your watch by it — an MIT program called Mac Hack VI, written by Richard Greenblatt, sat down against Dreyfus himself and beat him, to the unconcealed joy of the artificial intelligence community, whose in-house bulletin reported the result under a headline I will paraphrase as: the ten-year-old can beat the machine, but the machine can beat Dreyfus. Both sides claimed vindication. Both were premature by decades. It is the oldest rhythm in this field — overclaim, ridicule, quiet progress, ambush — and Chapter Four has already shown you what it did to the neural-network people during the same years.

The quiet progress happened in leagues. From 1970 the computer people held annual tournaments in which programs played *each other* — machine

versus machine, humans reduced to porters and seconds — and the first world computer championship, in Stockholm in 1974, was won not by the Americans but by a Soviet program, Kaissa; a few years earlier a program from a Moscow physics institute had beaten Stanford's best in a match played by telegraph over nine months, the Cold War conducted at one move per day. It was a Soviet mathematician, Alexander Kronrod, defending his laboratory's scandalous use of scarce computer time on games, who coined the field's best justification: chess, he said, is the *Drosophila* of artificial intelligence — the fruit fly, the small fast-breeding organism on which you work out the genetics before you touch anything that matters. The whole of this chapter is contained in that phrase. Nobody funds fruit flies for love of flies.

And the hardware story produced, en route, one episode so preposterously prophetic that I must slow down for it. Ken Thompson of Bell Labs — the same Thompson who co-created Unix, dry, brilliant, allergic to fuss — decided in the late 1970s that the way to make a chess machine strong was to stop being clever and start being *fast*: he and Joe Condon built Belle, a machine whose circuits were chess — special-purpose hardware that did nothing else, examined positions by the hundred thousand per second, and duly won the world computer championship and became the first machine granted the rank of US master. In 1982 Thompson packed Belle for a friendly exhibition in Moscow, and the United States government seized it at Kennedy airport as an unlicensed export of strategic technology to the Soviet Union. A chess computer. Arms control, applied to a game machine. The computing world laughed — colleagues suggested the only way Belle could threaten Moscow was if you dropped it on someone — and the laughter lasted, on and off, for forty years, right up to the day it stopped, because as I write this the export of game-derived chips is discussed at summit level between great powers, in earnest, precisely as if they were weapons. The customs men at Kennedy were not wrong. They were forty years early, and about the wrong box.

Belle's philosophy — special hardware, brute speed — went to graduate students at Carnegie Mellon, whose machine Deep Thought was hired, students and all, by IBM and renamed Deep Blue, and so we arrive at New York in May 1997, the thirty-fifth floor of the Equitable Center, and the most famous parlour game since Schönbrunn. Garry Kasparov, world champion, the strongest human player in the recorded history of the game, had beaten an earlier Deep Blue in Philadelphia the year before and had agreed, with the confidence of a man who had never lost a match to anything, to a rematch against a version examining two hundred million positions per second. What happened next was not really a chess story; it was a ghost story, and the ghost was Kempelen's. In the first game, which Kasparov won, the machine at one point played a move so pointless that it stuck in his mind — and one of Deep Blue's fathers, Murray Campbell, suggested years afterwards that the move was a *bug*: the program, failing to settle on an answer, had played more or less at random. Kasparov's camp, analysing overnight, drew the opposite conclusion: a move that inexplicable must conceal unfathomable depth. Poe's fallacy, running in reverse, in the mind of the best chess player on Earth. In the second game the machine declined a pawn-grab any computer was supposed to take, played with a positional patience the champion regarded as a human signature, and Kasparov — rattled, seeing hidden grandmasters the way the Viennese had seen clockwork — resigned. The position he resigned, analysts demonstrated within a day, was a draw; he had abandoned a game he could have saved by perpetual check, the first resignation of a drawn position in his life. He never recovered. He accused IBM, in everything but the plain words, of having a man in the box — reaching at one point for football's most famous handball, the hand of God — and demanded the machine's logs; IBM gave him extracts and stonewalled the rest. In the sixth and final game, needing only composure, the champion played a known losing line in the opening — a line whose refutation requires a piece sacrifice no cautious *human* would risk, and which the machine,

knowing no caution, played instantly — and was destroyed in nineteen moves, in under an hour, in front of the cameras. Deep Blue won the match. IBM, having got what it came for — the share price behaved as share prices near miracles do — declined a rematch, dismantled the machine, and sent one of its two towers to the Smithsonian. Match, machines. And yet: two hundred million positions a second is not thinking; it is counting, heroically. Deep Blue understood chess the way an avalanche understands skiing, and everyone in the field knew it. The fruit-fly experiment had produced a magnificent result and no genetics. The interesting animal was already breeding elsewhere, in a game most scientists considered beneath them, played with dice, in the back pages of the same IBM research reports.



Backgammon resists brute force for the most elegant of reasons: dice. Every turn multiplies the future by the twenty-one ways two dice can fall, so the tree of possibilities becomes a jungle within three moves, and Shannon-style deep search suffocates. In the early 1990s an IBM researcher named Gerald Tesauro tried something else. His program, TD-Gammon, contained a small artificial neural network — a descendant of the frog-twitch lineage you followed in Chapter Four — and it was taught backgammon by nobody at all. It played against itself, Samuel-fashion, hundreds of thousands of games, adjusting its own judgement after every one by a method called temporal-difference learning, which is a mathematician's way of saying: notice how things turned out, and revise your opinions of the positions that led there. From random flailing it climbed, game by game, to the level of the best players in the world. And then came the detail that matters: the world's best players started copying *it*. TD-Gammon evaluated certain standard opening rolls in ways that contradicted a century of received wisdom, the professionals tested its heresies, and the heresies won; expert practice shifted; the books were rewritten. For

the first time — quietly, in a game the newspapers ignored — a machine had not merely matched human mastery but *taught* the masters something. No grandmaster copied Deep Blue's style, because Deep Blue had none. Everyone copied TD-Gammon's. Counting had met learning, and learning was the one with something to say.

Which brings us to Go, and Go is where the whole nineteen-fold story pays off, so let us take it slowly. Go is older than chess by a millennium or two, simpler in its rules than draughts — surround territory, capture stones — and vaster in its possibilities than both by margins that embarrass arithmetic. The number of legal positions on the nineteen-by-nineteen board was finally computed, in 2016, at around two times ten to the power one hundred and seventy; the number of atoms in the observable universe is about ten to the eighty, a comparison not close enough to be called unfair — the atoms simply are not in the game. Brute force, the Deep Blue road, was therefore closed — not difficult, *closed* — and for two decades after Kasparov fell, the best Go programs lost to talented children, and Go players permitted themselves a certain serenity. Their game, they explained, required intuition: the strong player does not calculate the jungle, he *looks* at the board and the good moves announce themselves, by a trained perception nobody could articulate, let alone program. They were right about the perception. They were wrong about the "nobody could program it."

The company that did it was called DeepMind, founded in London in 2010 by, among others, Demis Hassabis — and pause on him, because his biography is this chapter in miniature. Chess prodigy at four; master strength in childhood; and at seventeen, professionally employed designing *video games*, co-writer of a hit called Theme Park, before going off to Cambridge and then into computational neuroscience. Games all the way down. His new company's first public triumph, in 2013, was a neural network that taught itself, from raw pixels and the score alone, to play the old Atari arcade games — and on Breakout, the one where you knock a

hole through a wall of bricks with a rebounding ball, the network discovered, entirely unaided, the trick every misspent teenage afternoon had discovered in 1976: tunnel through one side, get the ball trapped behind the wall, and let physics collect the points while you rest. Google saw that demonstration and bought the company. Keep the accounting straight here, because history was rhyming so hard it practically scanned: an artificial mind proved its worth to the largest information company on Earth *by playing an arcade game from the 1970s* — and what the 1970s arcade had already done for artificial minds is the other half of this chapter.

AlphaGo, DeepMind's Go machine, combined deep neural networks — one proposing moves, one judging positions, both trained first on human master games and then, decisively, on millions of games *against itself*: Samuel's trick, Tesauro's trick, now at planetary scale — with a search method that sampled the jungle rather than felling it. In October 2015, in secret, it played the European champion, Fan Hui, and won five games to nil: the first time a professional had lost to a machine on even terms. And in March 2016, in a ballroom of the Four Seasons hotel in Seoul, it sat down — its human operator sat down, rather; the machine itself was a data centre elsewhere, the Turk's cabinet grown to the size of a building and the man and the machine having at last fully swapped places — against Lee Sedol, winner of eighteen international titles, the finest player of his generation. Lee had studied the Fan Hui games and announced, reasonably, that he expected to win five-nil or at worst four-one. Two hundred million people, by Google's count, watched some part of what followed, an audience no board game had drawn in the whole history of board games.

Lee lost the first game and was visibly shaken. Then, in the second game, came the moment this chapter has been travelling towards since Schönbrunn: move thirty-seven. AlphaGo played a shoulder hit on the fifth line — never mind the technicalities; what matters is that it was a move essentially forbidden by centuries of accumulated master wisdom, so contrary to human instinct that the machine itself calculated a human

would play it about one time in ten thousand. The commentators assumed, live on air, that it was a mistake, possibly an operator's misclick. Lee Sedol left the room, was gone a quarter of an hour, and on returning sat looking at the move for minutes more before replying. Fan Hui, who after his defeat had joined DeepMind as a tester, said of move thirty-seven: it's not a human move; I've never seen a human play this move; so beautiful. *Beautiful* — from the man it had humiliated. Fifty moves later its quiet influence decided the game. Move thirty-seven was not a calculation dressed as intuition, the Deep Blue counterfeit. It was intuition — trained perception, the very faculty the Go world had declared unprogrammable — acquired the same way the masters acquired it: by playing, endlessly, against the strongest available opponent, which happened, in the machine's case, to be itself.

Humanity got one moment back, and it is right to record it. In the fourth game, long past hope at three-nil down, Lee Sedol thought for a very long time and played a wedge at move seventy-eight that the machine had assessed, by lovely symmetry, at about one in ten thousand — Korean commentators immediately christened it the divine move — and AlphaGo, thrown into futures it had never rehearsed, unravelled, blundered like an amateur, and resigned. One game to four. Lee said afterwards that he would not exchange that single victory for anything, and for one evening in Seoul a win at a board game by a member of our species over a machine was front-page news on six continents, which tells you everything about how far we had come from Schönbrunn, where the fraud had run the other way. Three years later Lee Sedol retired from professional play, and gave his reason in a sentence I ask you to read twice, because it is the sound of a door closing on three thousand years: even if I become the number one, he said, there is an entity that cannot be defeated. And the entity in question had already been superseded — a successor, AlphaGo Zero, had been given nothing but the rules, no human games at all, and in three days of playing against itself from pure ignorance it had

climbed past every master who ever lived and beaten the version that beat Lee by one hundred games to nil. Three days. The whole climb — from beginner's flailing, through every plateau that human tradition took centuries to cross, to somewhere beyond our sight — in three days, alone, by play. Evolution found the same trick long ago, which is why kittens pounce on string and wolf cubs wrestle: play is the safe rehearsal of the dangerous world, the training of judgement where mistakes are cheap. It is not the opposite of serious work. It is how nervous systems do their serious work, and in 2017 we finally learned that the hard way, from the machine.



And now I have a promise to keep. The last chapter left a small band of unfashionable believers in the neural network holding an idea that worked in principle and starved in practice, waiting — though they did not know it — for computing power of a very particular kind, from, I said, an unlikely source. Here is the source. It is a pinball machine. Near enough.

In 1962, on a PDP-1 computer that the manufacturer had donated to MIT more or less to see what the local youth would do with it, a young programmer named Steve Russell and a loose confederacy of friends — self-described members of the Hingham Institute Study Group on Space Warfare, an institution that existed chiefly on a sofa — wrote a program in which two spaceships duelled with torpedoes around a star whose gravity one of them had insisted on making real, while another, offended by the fake random stars, programmed in the actual night sky from a star catalogue, because if a thing is worth doing at all it is worth doing properly. They called it *Spacewar!*, exclamation mark included. It served no purpose whatever. It spread to practically every research computer in America — the manufacturer eventually shipped it as test software, so that the final check that your half-million-dollar computer worked was, officially, a game — and among the students it captivated, at the University of Utah, was an engineering undergraduate who paid his way working the

midway games at an amusement park. Watch this man. His name is Nolan Bushnell, and he is a bridge with a person attached: on one side of him, the research computing of the sixties; on the other, the carnival — the barker's patter, the coin box, the deep science of separating a teenager from twenty-five cents. In 1971 he built a coin-operated Spacewar clone; it flopped, being too clever for the tavern. In 1972 he adjusted the cleverness, founded a company, and — here is a loop even I would hesitate to invent, and I remind you I did not — this lifelong player of Go named his company after the moment in Go when a stone stands one move from capture, the game's equivalent of check. He called it Atari. The company named for the oldest game's moment of mortal threat would sire the industry whose chips, half a century later, put the oldest game itself in atari. If you wrote it in a novel they would make you take it out.

Atari's first product was written as a training exercise by its first engineer, Al Alcorn: bat, ball, two knobs, a scoreboard, and the now-sacred instruction, the entire design document of the modern age, printed on the cabinet: AVOID MISSING BALL FOR HIGH SCORE. They installed the prototype of Pong in Andy Capp's Tavern in Sunnyvale, California, and within a fortnight it broke down, and Alcorn drove over to fix it, and found the fault: the coin box was jammed. Too full of quarters. Engineers dream all their lives of a failure mode like that. (Honesty compels a parenthesis: a television engineer named Ralph Baer had patented much of the idea first, and his lawyers compelled a settlement, and the games industry thereby established in its infancy the tradition of litigation it has maintained with undimmed enthusiasm ever since.) The quarters became a flood, the flood became an industry, and in 1978 the industry produced its Beethoven's Fifth: Space Invaders, from Japan, designed by Tomohiro Nishikado — a game so consuming that legend insists Japan ran short of hundred-yen coins, a legend too tidy to be true, which has never once prevented its repetition, including, you will observe, here. Space Invaders contains my favourite accident in all

engineering. Nishikado's hardware was too feeble to draw all the invaders quickly, so the game ran slowly at the start; as you shot invaders there was less to draw, so the survivors sped up, and the march that begins as a lumber and ends as a sprint — the accelerating heartbeat that made the game terrifying and the terror that made it profitable — was a *bug*. The processor's weakness became the game's soul. Hold on to that, too: this entire industry is about to be built on the discovery that drawing things fast is the whole problem, and the whole prize.

Because now begins the arms race, and it is the strangest arms race in industrial history, in that its munitions were pictures and its battlefield was the inside of a teenager's skull. Every year, for thirty years, in arcades and then living rooms, the game that *looked* better took the quarters, and so every year the industry demanded chips that could draw more, faster: more sprites, then — the decisive turn, in the nineties — three-dimensional worlds built from thousands upon thousands of tiny triangles, each one shaded, textured, lit, and splashed onto the screen sixty times a second. And drawing triangles, it turns out, has a peculiar mathematical personality. Each corner of each triangle must be pushed through the same little sequence of multiplications; each pixel must be coloured by the same small recipe; and — here is the crux, so I will lean on it — none of these millions of tiny calculations needs to wait for any of the others. They can all be done *at once*. The computers of Chapter Two's lineage, the noble line of Babbage, do one thing at a time, very fast, in sequence, like a single brilliant clerk. Drawing explosions calls for the opposite creature: ten thousand very stupid clerks, all doing the same trivial sum simultaneously on different scraps of the picture. So the games industry, caring for nothing but wetter fire and shinier chrome, spent its rivers of pocket money breeding exactly that creature — the massively parallel arithmetic engine — and sold it, at consumer prices, in the tens of millions, to children.

The firm that came to own this trade was founded in 1993 by three engineers around a table at a Denny's diner in San Jose — Jensen Huang,

who had scrubbed dishes at a Denny's as a teenager and evidently regarded the establishment as lucky, Chris Malachowsky, and Curtis Priem. (The Introduction, you may recall, spoke of engineers whose only ambition was to make video-game explosions look wetter and shinier; here they are, and at the time this was, give or take some diplomatic paraphrase, the actual business plan.) They called the company NVIDIA, their first chip was a commercial failure, they came within a payroll or two of extinction — Huang took to opening meetings with the cheerful observation that the company was thirty days from going out of business, a sentence he kept using long after it stopped being true — and they clawed back with chips that did one thing supremely: triangles, fast. In 1999 they shipped a card called the GeForce 256 and coined, for marketing purposes, a term of art: the Graphics Processing Unit. The GPU. The initials were pure salesmanship. Remember, as you watch those three letters take over the world, that they were invented to sell Christmas presents.

Then, around 2001, the graphics chips crossed a threshold almost nobody outside the trade noticed: they became *programmable*. To let game artists concoct ever-fancier fire and water, the pixel machinery was opened up so that developers could load their own little recipes — shaders, in the jargon — and at that moment a few impecunious scientists, squinting at the specifications, realized what was now sitting inside the family Christmas present: a supercomputer. A cramped, eccentric supercomputer that spoke only the language of pictures — but parallel arithmetic is parallel arithmetic, and the multiplication of enormous grids of numbers, which is the daily bread of physics, of chemistry, and (Chapter Four) of neural networks, is exactly the kind of work ten thousand stupid clerks do best. So the scientists began to smuggle: they told the chip it was painting a picture when it was actually simulating a molecule, painting their mathematics as textures and reading off the answers as colours. Kempelen's cabinet, turned inside out one more time — serious thought, hiding inside a toy, pretending to be entertainment in order to be allowed to exist.

NVIDIA noticed the smugglers, and here credit where due: instead of prosecuting the metaphor, the company legalized it. It hired the Stanford student whose thesis had made the smuggling systematic, and in 2007 shipped CUDA — a scheme by which any programmer could address the games chip in plain arithmetic, no disguises, no pictures — and, crucially and expensively, built the capability into every chip it sold, gaming cards included, spending silicon on scientists who did not yet exist as a market. Wall Street grumbled for years about the waste. The graphics card in a teenager's bedroom in 2008 was, whether the teenager knew it or not, a scientific instrument awaiting instructions. And in 2012 the instructions arrived. In Toronto, in the laboratory of Geoffrey Hinton — one of the stubborn believers whose thirty years in the wilderness you have already read about, so I need not repeat the sermon — a graduate student named Alex Krizhevsky trained a deep neural network on a million labelled photographs for the great annual image-recognition contest, and he trained it at home, in his bedroom, on two NVIDIA GTX 580s: gaming cards, bought like any other Christmas present, about a thousand dollars the pair. The training took rather less than a week — the same work would have taken a conventional processor months into years, which is to say would simply never have been done — and the network, AlexNet, did not so much win the 2012 contest as end it: an error rate of fifteen percent against the runner-up's twenty-six, a margin that in that field was not a victory but a geological event. Within eighteen months everybody — every lab, every language, every task, vision and speech and translation and, in the fullness of time, the machine now writing books about itself — was training neural networks on games chips, because there was nothing else on Earth that could do the arithmetic at any price a scientist could pay. The winters of Chapter Four ended not because anyone built the believers their machine, but because the believers looked up and found that the amusement arcade, for its own disreputable reasons, had spent thirty years

building it for them. The explosion-drawing engine was the mind-growing engine. The toy was the loom.

You know the rest, because you are living in it. The company from the Denny's booth became, for a while at least, the most valuable company on Earth — worth more than the oil majors, the banks, and the whole respectable establishment that had spent decades regarding games as beneath serious attention. Its chips are now, ounce for ounce, among the most precious objects industrial civilization produces: inventoried like bullion, escorted like bullion, restricted like weapons — and if you want to know how seriously to take the customs men who impounded Ken Thompson's chess machine in 1982, the answer is now on the front pages several times a week. And the pattern, please notice, ran start to finish without a single planner. Nobody at MIT in 1962, at Andy Capp's Tavern in 1972, at the Denny's in 1993, was building infrastructure for artificial minds. Each link wanted only the next quarter — in, for most of the chain, the literal coin-box sense. I once wrote a book organized around the way change ricochets off the incidental, and called the phenomenon the pinball effect, and I ask you to admire how thoroughly this chapter has abused my hospitality: the pinball effect, delivered *by the pinball machine's own descendants*. The amusement business built the engine of the next civilization by accident, twenty-five cents at a time.



So: should we have seen it coming? Yes. And we have no excuse, because play has been running this operation on us from the beginning, and the file is thick.

Hero of Alexandria, in the first century, built a sphere that spun by the power of steam, and his contemporaries used steam and its cousin pneumatics for temple wonders and party tricks — doors that opened by themselves when the altar fire was lit, birds that sang, wine that poured — and then everyone went back to work, the work being done by slaves, and

the steam engine waited seventeen centuries for someone to need it. David Brewster invented the kaleidoscope in 1816 as a by-product of serious optics, and by his own estimate — handle inventors' estimates, like sons' anecdotes, with tongs — two hundred thousand were sold in London and Paris inside three months: a fad, and the lens-grinding trade it fed was the trade that supplied the century's microscopes and cameras. Michael Faraday in 1821 made a wire circle endlessly around a magnet — the first conversion of electricity into continuous motion, the ancestor of every motor in your house — and polite science received it as an ingenious *toy*; the anecdote in which he answers a scoffing statesman with "one day, sir, you may tax it" is, I regret to report, almost certainly a later invention, but the scoffing was real, and the toy runs the world. The pattern is so reliable you could invest by it (some people, it will not have escaped you, now do): the significant technology arrives disguised as an amusement, because an amusement is allowed to be useless, and only what is allowed to be useless is free to become whatever it wants. Evolution, as I said, discovered this long before we did. Play is the R&D department of the vertebrates — the department where failure is survivable, which is the only department where genuine novelty can afford to happen. The kitten stalking a ball of wool is running combat simulations on a training budget. So was Turing, executing his imaginary chess machine with a pencil. So was the network that learned Breakout. So, if we are honest, was the whole species, every time it wasted an afternoon on a board of nineteen lines and a bag of stones. Three thousand years of apparently idle play turned out to be a specification, written in advance, for the first thing we would want a mind-made-of-machinery to do — and thirty years of allegedly misspent youth in the world's arcades turned out to be the capital campaign that paid for the machinery. Your aunt — the one the machine is learning to imitate — always said those games would amount to nothing. It is one of the very few things she was wrong about, and I look forward, in a later chapter, to hearing the machine tell her so itself.

But before it can say anything at all, someone has to work out what saying *is* — what messages are made of, why some sequences of letters carry meaning and others carry noise, and how you would even begin to measure the difference. That someone is the man I left at the end of the bar earlier in this chapter, the one with the unicycle and the mechanical mouse, the man who took the guessing games we play — hangman, twenty questions, finish-this-sentence — and found inside them the mathematics of all communication whatsoever. His story begins, of all places, on a chain of French hilltops, with an eighteenth-century engineer waving wooden arms at the horizon.

Which is the next chapter. Bring a scarf. The hilltops are windy.

CHAPTER SIX

Getting the Message

*We are in great haste to construct a magnetic telegraph from
Maine to Texas; but Maine and Texas, it may be, have nothing
important to communicate.*

— HENRY DAVID THOREAU, *WALDEN* (1854)

For about two years in the mid-1830s, an odd thing kept happening on the French government telegraph line that ran south-west from Paris to Bordeaux. Every so often, in among the official dispatches, there would be a mistake. Just one wrong symbol, appearing somewhere after the city of Tours, followed a little later by the special signal that meant *cancel that, error, disregard* — at which point the mistake was duly struck from the record and everybody moved on, because operators were human and errors were routine and the error-and-cancel procedure existed precisely so that nobody need think twice about it. Nobody did think twice about it. That was the idea.

Because the mistake was the message. On a hill outside Bordeaux, a man with a telescope was watching for it. The wrong symbol told him which way government bond prices had moved in Paris — days before the news could arrive by mail coach, which was the fastest thing the public had. And the men paying for that telescope, and for the bribed operator at Tours who injected the "error" into passing government traffic, and for the accomplice in Paris who watched the market and signalled the day's movement down to Tours by the colour of the wrapping on otherwise innocent parcels — those men were two Bordeaux bankers, twin brothers, named François and Louis Blanc. For two years they beat the market with information nobody else in Bordeaux could have, using a state communications network that private citizens were forbidden to touch, by hiding their signal inside the system's own tolerance for noise. It was insider trading by telegraph. It was also, as near as anyone can tell, the first hack of a data network in the history of the world, complete with what a modern security man would recognize instantly as a covert channel and a corrupted node.

The scheme came apart in 1836 in the usual way — one of the bribed operators fell ill, tried to recruit a successor, and the successor talked — and then came the truly delicious part: at the trial, the court discovered that no law existed against any of it. Nobody, drafting the statutes of France, had imagined that information travelling at the speed of sight was a thing that could be stolen. The brothers walked free, and the law scrambled along behind the technology, as it always does, and has been scrambling ever since. Alexandre Dumas, who never let a good crime go to waste, put a bribed semaphore operator into *The Count of Monte Cristo* a few years later and used him to ruin a banker, which tells you the story was common gossip. And François Blanc — I promise I am not making this up — took his talent for privileged odds off to the gaming tables, ran the casino at Bad Homburg, and ended up building Monte Carlo, where he grew spectacularly rich on the one form of gambling in which the house

knows the probabilities and you don't. The first man to hack a network spent the rest of his life monetizing other people's uncertainty. Hold that thought. This entire chapter is about monetizing uncertainty, and by the end of it you will know what the talking machine does all day.

But first: what exactly had the Blancs hacked? What was a telegraph doing on French hilltops in 1834, decades before the electric one? For that we have to go back to the Revolution, and to a defrocked provincial cleric with a telescope and a very determined family.



Claude Chappe had been comfortably set up as a young *abbé* with church income and a taste for physics experiments, until the Revolution abolished the income, at which point — necessity, invention, you know the proverb — he turned his full attention to a problem he had been playing with for years: how do you send a message faster than a horse? His first public experiment, on the 2nd of March 1791, sent a sentence about ten miles between the towns of Brûlon and Parcé using synchronized clocks and a painted panel, before dignitaries who signed affidavits to say they'd seen it. The sentence, dictated on the spot by the local worthies to prove nothing had been pre-arranged, was: *if you succeed, you will soon bask in glory*. Note it well, because it is possibly the least accurate message ever transmitted by telegraph, and telegraphs have carried some whoppers.

The machine that eventually worked was beautifully simple: a tall mast on a tower, carrying a pivoting wooden beam, with a smaller pivoting arm at each end of the beam, the whole thing worked from below by ropes and cranks like an enormous marionette. Set the beam and the arms at various angles and you could form ninety-two distinct, unambiguous shapes, readable through a telescope from the next tower, five or ten miles away, where an operator copied the shape for the tower after that. Chappe wanted to call it the *tachygraphe* — the fast writer — until a classically educated friend in government, Miot de Mérito, suggested that the interesting thing

about it wasn't that it wrote fast but that it wrote *far*, and gave it the name we still use for its successors: telegraph. Getting the state to care was the hard part, and here it helped enormously that another Chappe brother sat as a deputy in the Legislative Assembly, because in revolutionary Paris an unexplained signalling machine was a good way to get killed: twice, early on, patriotic mobs burned Chappe's apparatus on the entirely reasonable suspicion that anyone sending coded messages over the heads of the people was talking to royalists. The Revolution wanted speed and feared secrecy, and the telegraph was both. It has been both ever since.

What sold it was the war. By 1794 France was fighting most of Europe, and the Committee of Public Safety approved a chain of towers from Paris to Lille, a hundred and forty miles towards the fighting on the northern frontier. That August it delivered the news that the town of Condé-sur-l'Escaut had been retaken from the Austrians — delivered it to the Convention in Paris within the hour. The Convention was so pleased it renamed the town Nord-Libre on the spot and telegraphed back that the Army of the North had deserved well of the fatherland, and the whole exchange — event, news, response, all in an afternoon, across a distance that had been three days by courier since the Romans — was, quietly, the end of one world and the start of another. Distance had been a fact. It was now a cost. Costs come down.

The network grew for sixty years: under Napoleon, who used it to move armies and, the day after his coup in 1799, to inform France that "Paris is quiet and the good citizens are content" (which was at least half true); under the restored kings; under the July Monarchy that the Blanc brothers embarrassed; until at its height five hundred and thirty-odd towers webbed France from Calais to Toulon, and other nations had copied the idea. Under good conditions a signal could flicker from Paris to Lille in a couple of minutes. There were only two problems. The first was night. The second was weather. A state-of-the-art national data network, and it went down every evening at dusk and every time it rained — which is worth

remembering next time your connection drops and you feel civilization tottering.

Now to the part of the story that matters most for our purposes, and it is not the towers; it is the codebook. Ninety-two shapes are not many, and spelling everything out letter by letter would have been unbearably slow. So the Chappes did something clever: they used the signals in pairs, the first to name a page in a codebook, the second to name a numbered entry on that page — ninety-two times ninety-two, eight and a half thousand words and stock phrases, everything a government routinely needs to say, pre-packaged. Two flaps of the arms and you had transmitted an entire military phrase. Notice what is going on there, because it took me years to notice it myself: the message hurtling across France was not language. It was a pair of *pointers* into a shared body of knowledge sitting in Paris and at the far terminal. The less surprising your message — the more it resembled what the system already expected people to say — the cheaper it was to send. Predictability is compressibility. A twentieth-century engineer will build a whole science on that sentence, and the Chappe brothers were living it in 1794.

And one more thing about those towers. The operators in the middle of the line — poorly paid men who watched the neighbouring tower through a telescope and copied what they saw, on pain of fines — did not have the codebook. Only the terminals did. Which means that for sixty years, the French national communications system was a chain of men each sitting in a box, receiving symbols he could not understand, consulting a fixed set of rules, and passing symbols on — while the *system as a whole* transmitted meaning perfectly well from end to end. Each man understood nothing; the line understood everything, if a line can be said to understand, and whether it can is precisely the question. File that image somewhere safe. Philosophy will reinvent it in 1980, as an argument about whether a machine that shuffles symbols it doesn't understand can be said

to think, and the row it causes will still be going when the machines start shuffling back.

Claude Chappe did not live to see most of this. The glory his brother promised curdled quickly: rivals and cranks hounded him with claims that the telegraph had been stolen — from them, from the ancients, from anyone — and Chappe, exhausted, unwell, and sunk in depression, went into the courtyard of the telegraph administration in Paris one January day in 1805 and threw himself down the well. He was forty-one. His brothers put a telegraph tower on his tomb, arms frozen in the signal for *rest*. The priority dispute, you will have noticed, is a recurring character in this book, and it kills more inventors than poverty does. Keep an eye on it; it appears again within a few pages, aboard a ship in the mid-Atlantic.



In the winter of 1825 an ambitious American painter named Samuel Finley Breese Morse was in Washington, doing rather well: he had the commission of the season, a full-length portrait of the Marquis de Lafayette, then making his triumphal tour of the republic he'd helped midwife. On the 11th of February a letter arrived from Morse's father in New Haven, written days earlier: Samuel's wife Lucretia, three weeks after the birth of their third child, was dead. He dropped everything and rode for home. It took him the better part of a week. By the time he got there she was already buried. He had not known she was even ill.

The story goes — and I give it to you as the story goes, because it was polished smooth by Morse's Victorian biographers and has been polished smoother since — that in his grief Morse resolved that no one should ever again suffer news arriving too late to matter, and so invented the telegraph. The tidy version is too tidy: the man went on painting for years, and ran for mayor of New York on a platform so poisonously anti-Catholic and anti-immigrant that one hesitates to bring it up in polite company (Morse was, let us be candid, a bigot of energy and conviction; history does not

distribute its gifts by merit). But the wound was real, and wounds have a way of setting a compass without telling the owner. When the idea finally arrived, it arrived — as these things so often do in our story — on a boat, in conversation, sideways.

October 1832, the packet ship *Sully*, Le Havre to New York, five weeks at sea and nothing to do but talk. A Boston doctor named Charles Thomas Jackson entertained the dinner table with the latest European parlour physics: the electromagnet, and the curious fact that electricity passes through any length of wire essentially instantly. Morse, forty-one years old and by his own lights a failure, seized on it with a convert's fervour. If the presence of electricity can be made *visible* at the far end of a wire, he said, I see no reason why intelligence may not be transmitted instantaneously by electricity. He spent the rest of the voyage filling a sketchbook, and the rest of the decade — broke, obsessed, living on his students' fees as an art professor — making the sketches work. (Dr Jackson, for his part, later claimed he had really invented the telegraph on that ship. He also claimed he had really discovered ether anaesthesia, and really invented guncotton, each time some years after somebody else had done the work; he died in an asylum, still owed, in his own accounting, most of the century's credit. I told you the priority dispute would be back.)

Two things stood between Morse's sketchbook and the world, and their names were Alfred Vail and money. Vail was the mechanically gifted son of a New Jersey ironmaster; he saw Morse demonstrate the crude prototype in 1837, bought in for a share of the rights, and in the machine shops of his father's Speedwell Ironworks turned a professor's contraption of canvas-stretcher and homemade battery into an instrument that clicked and worked and could be manufactured. How much of "Morse's" mature system is actually Vail's has kept historians happily quarrelling for a century and a half — Morse, true to the pattern of this chapter, spent his later years litigating and lobbying to make sure no other name stuck to the thing — but on one point the evidence tells a lovely story either way.

Morse's original plan for sending language, you see, was a numbered dictionary: every English word assigned a number, the numbers sent as pulses, a codebook at each end. It was, detail for detail, the Chappe system over again — pointers into a shared book — and it was a dead end, because English is not eight thousand stock phrases, and looking up every word is slow. The system that replaced it did something subtler: it spelled, letter by letter, but it made the letters *cost* differently. E, the workhorse of English, became a single dot — the cheapest signal possible. T, a single dash. Q and Z and the other aristocratic idlers of the alphabet got long, expensive combinations. The code was a wager on the statistics of the language, shaped so that the average message would be as short as the language allowed.

And where did the statistics come from? Not from any survey, census or learned society. The story — it comes down through the Vail side, and I believe it because it is too practical to be invented — is that the type case of the local printing office in Morristown was consulted, and the answer simply read off the compositors' boxes. A printer's case is a wooden tray of compartments, one per letter, and the compartments are not the same size, because three centuries of setting English by hand had taught the trade exactly how many of each letter a page consumes: a font of type shipped with hundreds of e's for every j, and the *e* box was the biggest in the case, sitting under the compositor's hand where the reach was shortest. (The capitals lived in the case racked above — the *upper* case; yes, that's where the words come from.) Nobody involved thought of this as mathematics. It was stock control. But it was a frequency table of the English language, compiled empirically by generations of printers who would have been astonished to learn they were doing statistics — and Morse code is that table, cast into dots and dashes. A century later, when information theory existed and somebody thought to check, the old code turned out to be within shouting distance of the theoretical optimum. The compositors had known. Language has a shape, the shape can be counted, and if you count it

you can compress it. Remember the sweet peas of Chapter Three, where we learned that averages have laws? Here is the same discovery arriving by the back stairs: *letters* have laws. The alphabet, taken in bulk, is as predictable as artillery.

The rest is the famous part. Congress, after years of Morse's lobbying, voted thirty thousand dollars in 1843 to string a demonstration wire from Washington to Baltimore — and very nearly attached an amendment, offered in a spirit of mockery, to spend half the money investigating mesmerism, the argument being that one invisible influence acting at a distance was as good as another. The bill passed by six votes, with seventy congressmen prudently abstaining rather than be caught on the record either way about the future — there is a whole book to be written about that seventy. On the 24th of May 1844, in the chamber of the Supreme Court, Morse tapped out a verse chosen by Annie Ellsworth, the patent commissioner's young daughter: *What hath God wrought* — Numbers, chapter twenty-three, verse twenty-three, and rather less a boast than we remember it, since in context it is an exclamation of astonishment, very nearly of alarm. Vail, in Baltimore, sent it straight back, and the message thus travelled eighty miles in a heartbeat and history divided itself neatly into before and after.



After, it went fast — you know it went fast, this is the pattern of the book — and I will spare you the mileage statistics, which were quoted breathlessly at the time in exactly the tone we now reserve for counting transistors. Within a decade wires laced the eastern United States and most of Europe; within three, you could cable from London to Bombay. Prices of goods, which had varied city by city because information travelled at the speed of a horse, converged into single national markets. News became a commodity with a wholesale trade: Reuters began as a man filling the telegraph gap between Aachen and Brussels with carrier pigeons.

Diplomacy, crime, cricket scores, the price of pork — everything that was information began to move at the new speed, and a historian of the thing has called the result, accurately, the Victorian internet.

And like our internet, it grew a culture, and the culture is where I want to linger for a moment, because buried in it is the chapter's secret subject. The operators — thousands of young men and, notably, young women, in an age that offered educated women almost no other technical employment — lived on the wire. In slack hours they gossiped across continents, played chess by wire, told jokes, courted. There were friendships between operators who never met, and marriages between operators who met only after proposing — there is an 1880 novel about exactly this, *Wired Love: A Romance of Dots and Dashes*, which reads today like a catalogue of internet-dating anxieties transposed a century early, up to and including the discovery that the witty stranger on the wire is not quite as advertised in person. Thomas Edison, who started as a telegraph operator and was half deaf, courted his second wife by teaching her Morse and proposed to her by tapping the question into her palm; she tapped back yes, which for my money is the most romantic use of channel capacity on record.

Here is the detail to keep. Every operator sent the same standardized code — that was the whole point of a code — and yet experienced operators could identify who was at the distant key within a few words, by rhythm alone: the weight of the dashes, the hesitation before a difficult word, the swing of it. They called it the operator's *fist*, and a fist was as recognizable as handwriting and as hard to fake. (In the Second World War, British intelligence identified German radio operators by fist, tracking individual men from posting to posting even when everything they sent was enciphered — reading the accent when the words were opaque.) Now think about what a fist is. The code specifies the letters absolutely; the fist is everything the code does *not* specify — and it turns out that this residue carries information too: not about the message, but about the sender. Style is measurable. Identity leaks through statistics. A pattern of tiny,

unconscious regularities is enough to tell the world who is speaking, whatever is being said. Chapters from now, when a machine reads a paragraph and makes a fair guess at its author — or writes one, and a second machine makes a fair guess that no human did — you will already know the principle: everyone who sends, signs.

One more piece of Victorian plumbing, briefly, because it forced the next question. In 1858, after two failed attempts and at ruinous cost, a cable was landed connecting Ireland to Newfoundland, and Queen Victoria exchanged telegrams with President Buchanan amid celebrations verging on the unhinged. Her ninety-eight-word message had taken sixteen hours to send: under two thousand miles of ocean, the crisp clicks that worked so well on land came through smeared into long ambiguous sighs. The cable's chief electrician, a confident surgeon-turned-engineer named Wildman Whitehouse, held that the cure for a faint signal was a bigger shout, and applied induction coils at some two thousand volts until, within a month, he had cooked the insulation somewhere under the Atlantic and the cable died. The man who had opposed him all along, a Glasgow professor named William Thomson, had the opposite instinct: don't shout — *listen better*. His mirror galvanometer hung a scrap of mirror on a silk thread and let a whisper of current swing a beam of light across a scale, so that a signal too weak to click a sounder could still be read by a patient human eye. Thomson's way won; the 1866 cable worked and kept working; he collected a knighthood and eventually the title Lord Kelvin, under which name he will turn up near the end of this book being wrong about aeroplanes — nobody is right twice in a row. But mark what had happened to the problem of communication. On an open hilltop it had been a problem of *codes*. Under an ocean it became a problem of *physics*: how fast can a channel be driven, how faint a signal pulled from the noise, what — precisely, numerically — are the limits? Nobody could answer, because the question wasn't yet mathematics. It was about to be, and the organization that made it mathematics was born, in the way of these things, from a different

gadget entirely: the one Mr Bell — who had once stared at Wheatstone's reconstruction of Kempelen's talking machine, as Chapter Five told you — patented in 1876, which turned the telegraph's clicks into the human voice and within two generations wired a continent for conversation.



The Bell System grew into the largest corporation on Earth, a regulated monopoly with a licence to collect a small toll on every conversation in America, and in 1925 it did something monopolies occasionally do with their embarrassment of margin: it endowed a research laboratory and told some very clever people, in effect, to go and think about whatever the telephone might one day need. Bell Telephone Laboratories would in time hand the world the transistor, the solar cell, the communications satellite and rather a lot of this book, and we shall be back inside its corridors shortly — on a unicycle. But the first thing its thinkers did was pick up the question the Atlantic cable had left lying on the seabed.

Two of them got hold of it properly. Harry Nyquist — Swedish farm boy, emigrated at eighteen, the sort of quietly monumental engineer whose name now lives in half a dozen theorems — showed in the 1920s that a channel's capacity to carry telegraph signals was fixed by its bandwidth in a precise, calculable way: so much bandwidth, so many distinguishable pulses per second, and no cleverness at the key could beat the number. His colleague Ralph Hartley went further and said the unsayable: that the "information" in a message could be measured, and measured without any reference whatever to what the message *meant* — simply as the logarithm of the number of possible messages the sender might have chosen. A choice among many possibilities carries much information; a choice among few carries little; a "choice" among one carries none. Note the scandal being born here, quietly, in a trade journal: meaning has just been escorted out of the building. The company's reasons were dollars — squeezing more conversations down one strand of copper

was worth millions, and you cannot optimize what you cannot measure — but the philosophical dynamite lay there in plain sight for twenty years, fuse unlit, waiting for someone with a match.

Before we meet the man with the match, however, I must take you to St Petersburg, twelve years before Bell Labs existed, to witness the strangest scene in this entire chapter: the first statistical model of human language being built by hand, out of a poem, in order to win a theological quarrel.



Andrei Andreyevich Markov was a Russian mathematician of ferocious ability and matching temper — a man who wrote letters of complaint the way other men breathe, who once demanded the Church excommunicate him too when it excommunicated Tolstoy (the Church, wearily, obliged). His enemy — every such man has one — was Pavel Nekrasov, a Moscow mathematician of mystical leanings, and their quarrel, incredibly, was about free will. The law of large numbers — Chapter Three again: the way individual chaos smooths into collective regularity, the drunkard's steps into the reliable average — had been proved only for *independent* events, coin tosses that know nothing of each other. Nekrasov seized on this: since social statistics (marriages, crimes, suicides) show exactly such smooth regularities, the underlying human acts must be independent like the coin tosses — uncaused by one another — and independence, said Nekrasov, is free will, mathematically demonstrated, and Christian doctrine is thereby underwritten by the St Petersburg census tables. Markov read this and experienced the emotion that produces mathematics. Rubbish, he said, in effect (his actual correspondence is less printable): the law of large numbers doesn't *require* independence, and I shall prove it — I shall exhibit a chain of events, each shackled to the one before, utterly *dependent*, and show that regularity emerges anyway, and Mr Nekrasov's syllogism will be dead, and with it his statistical theology.

He needed a chain of linked events, thousands of them, cheaply available for counting. He chose the first twenty thousand letters of *Eugene Onegin*, Pushkin's verse novel, the most beloved text in the Russian language, and he sat down — this is 1913, no machinery on Earth to help him — and counted, by hand, letter by letter, tallying how often a vowel followed a vowel, a vowel a consonant, and so on through the pairs. The letters of a poem are the perfect specimen of dependence: what comes next hangs on what came before; Pushkin is not tossing coins. And Markov showed that even so, the proportions settle, the averages converge, the law holds — dependence and all — and Nekrasov's free-will engine was scrap. As pure mathematics it founded the theory of what we now call Markov chains, which run today inside everything from page rankings to weather models to the pricing of the bonds the Blanc brothers once cheated on.

But look at what he did on the way, because it is the hinge of this book. To count those pairs, Markov had treated the masterpiece of Russian literature as a *statistical object* — had said, in effect: never mind Tatiana and Onegin, never mind the duel and the wasted love and the snow; regard this text as a stream of symbols, each occurring with a certain probability given its predecessor, and the stream has a structure, and the structure can be captured in a table of numbers. Given the letter before, his little table would tell you the odds on the letter after. It is a model — crude, hand-cranked, but genuine — of the *statistical texture of language*, built from a body of text. Every talking machine on Earth is a descendant of that table. The first language model was compiled with pen and paper by an angry atheist counting vowels in Pushkin to spite a mystic, and if you can find me a better example of the future arriving through the tradesman's entrance, wearing a disguise, on an errand of pure spite — I shall want to see the paperwork.



Now the man with the match. Claude Elwood Shannon was born in 1916 in Gaylord, Michigan, a small town where his father sold furniture and his mother ran the high school, and where the boy showed every symptom of the classic American tinkerer: model planes, radio sets, a barbed-wire fence pressed into service as a private telegraph line to a friend's farm half a mile off — young Claude signalling in Morse down the property lines of Michigan, a detail so on-the-nose that no novelist would dare use it. He worshipped Edison (a distant cousin, it later turned out, to Shannon's lifelong delight — both descended from the same colonial New Jersey stock), and he had the trait that marks everyone important in this chapter: an interest not in what messages *say* but in how they *move*.

At MIT he found his instrument. He worked as a graduate assistant on Vannevar Bush's differential analyser, the great room-sized mechanical brain of the age, and his particular chore was its tangled electrical control system — a hundred-odd relays, the clicking electromagnetic switches the telephone company used by the million to route calls. Relays are two-state devices: open or closed, on or off. And Shannon happened to be carrying, from a philosophy course back in Michigan, a piece of intellectual apparatus almost nobody else in engineering possessed: the algebra of George Boole, the self-taught son of a Lincoln cobbler who in 1854 had published *The Laws of Thought*, reducing logic itself to a calculus in which propositions are variables, true is 1, false is 0, and reasoning is arithmetic. Boole's algebra had sat for eighty years as a philosophers' curiosity, admired and useless. Shannon looked at the relays, looked at Boole, and saw that they were the same thing: a circuit is a proposition; switches in series are AND, switches in parallel are OR; the design of switching circuits, until that moment a black art of experienced tinkerers, is Boolean algebra made of brass and wire — and therefore any statement logic can express, a machine of switches can compute. His 1937 master's thesis said so, with examples, at the age of twenty-one. It has been called — the line is a Harvard psychologist's, but nobody has seriously contested the billing —

possibly the most important master's thesis of the century. It is why the machine on which you may be reading this sentence is, at bottom, several billion relays clicking out syllogisms. The axemaker's gift of Chapter One was language that could take the world apart and describe it piece by piece; here, in a student's homework, the taking-apart is handed over to the switches. Thought — some of thought, the dry sequential part — had become a substance you could manufacture.

Bell Labs, naturally, took him. The war put him onto cryptography and the transatlantic scrambled-voice link by which Roosevelt and Churchill conferred (and onto lunchtime conversations, in the Labs cafeteria in 1943, with a visiting English cryptanalyst named Alan Turing, where the two of them alarmed the neighbouring tables by discussing what else machines might one day do; Turing's own story is told in Chapter Five, where the games are kept). Cryptography, if you think about it, is information theory's evil twin — the study of how to make a message statistically indistinguishable from noise — and wrestling with secrecy all day sharpened exactly the question Nyquist and Hartley had left lying about: what *is* this stuff, this whatever-it-is, that channels carry, codes compress and ciphers hide? Shannon worked on the answer for eight years, mostly at night, mostly telling no one, with no deadline and no assignment — the monopoly's loose change at work — and in 1948 he published, in the company's house journal, a fifty-odd-page paper with the flattest title of the century: "A Mathematical Theory of Communication." There are perhaps three or four papers in history that create an entire science in one go, fully armed. This is one of them.

What does it say? First, with the opening sentence, it fences the ground: "The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point." Selected — mark the word. And then, immediately, the scandal made official: messages frequently have meaning, Shannon notes, but "these semantic aspects of communication are irrelevant to the engineering

problem." What matters is only that the actual message was *chosen from a set of possible messages*, and information is the measure of that choice: how much uncertainty, on the receiving end, does the message resolve? The unit — the word was suggested down the corridor by John Tukey, as a contraction of *binary digit* — is the bit: the amount of information in one fair coin toss, the answer to one yes-or-no question that could honestly have gone either way. Everything transmissible — the Ninth Symphony, the Book of Kells, your mother's voice saying your name — is so many bits, and the bit count depends not on profundity but on *improbability*. A message you already knew was coming carries nothing at all, however long it takes to say; which is why, Shannon might have added but was too polite, most after-dinner speeches contain no information whatever.

On this foundation he built, in one paper, the whole cathedral: a measure of the information a source produces per symbol; a measure of what a channel can carry per second; and the theorem joining them, which is the genuinely miraculous part. Everyone since Kelvin's cable had assumed that noise sets a floor under errors — a crackly line means garbled messages, and the only cures are shouting (see Whitehouse, above) or repetition. Shannon proved this is simply false. Every channel, however noisy, has a definite capacity, a number of bits per second; and below that capacity, messages can be encoded so cleverly — wrapped in exactly the right kind of calculated redundancy — that they arrive with as few errors as you please. *As few as you please*: not fewer errors at great cost, but essentially none, through the storm, provided you respect the speed limit and wrap the parcel properly. Above capacity, no ingenuity will save you; below it, perfection is available to anyone with the right code. It is why a photograph can cross the void from the far side of Saturn, whispered from a transmitter running on less power than a refrigerator bulb, and arrive without a freckle out of place. The noise doesn't win. It's the single most optimistic theorem I know.

For his measure of uncertainty Shannon needed a name, and the formula happened to be one physics already knew, from thermodynamics of all places. The story — Shannon told it himself, though he told it more than one way over the years, which is exactly what you'd expect from the man who measured how much of every story is padding — is that John von Neumann told him: call it *entropy*. First, because the mathematics really is entropy's mathematics. And second — the reason the story survives, polished or not — because "nobody knows what entropy really is, so in a debate you will always have the advantage." Entropy it became, and the joke came true, and the advantage has been enjoyed in seminar rooms ever since.



Now we come to the part of the theory that this book has been driving at since page one, and I want you to take it slowly, because when people say the talking machine is "just statistics" this is the statistics they are talking about, and they are right, and they have no idea what they are conceding.

A source of messages, says Shannon, has an entropy: the average genuine information per symbol, the residue of real unpredictability once all the pattern is accounted for. So — what is the entropy of *English*? If the twenty-six letters and the space were used independently and uniformly, each letter would carry about four and three-quarter bits. They are not so used, of course; Markov's ghost sees to that. E outnumbers Z (the printers knew); Q drags U along; TH, HE and AN do half the language's work; after *informatio* the next letter is not in serious doubt. Shannon worked the tables — single letters, pairs, triplets, whole words — and found the redundancy of ordinary English to be roughly fifty per cent. Half of everything you write is, in the engineer's cold sense, not information: it is scaffolding, predictable from the rest, removable and reconstructible. (He also noticed, in passing, the connection to that great Victorian information-game, the crossword: it is precisely because English is about

half redundant that words can be made to interlock in two dimensions — a language with no redundancy would make every random grid of letters a valid crossword, which sounds like heaven and would be hell, while a much less redundant language would support crosswords in three dimensions. This is the quality of Shannon's mind at play: a throwaway paragraph, and puzzle-setting becomes a branch of geometry.)

But the letter-tables only see a few characters back. What is the redundancy of English to a reader who brings *everything* — grammar, idiom, plot, gossip, the ways of detectives and of blondes? No table could hold that model. So in 1951 Shannon published one of the great cheap experiments of all time, "Prediction and Entropy of Printed English," whose apparatus consisted of a detective novel and a spouse. He had married, in 1949, Mary Elizabeth Moore — Betty — a numerical analyst at the Labs, one of the women employed (the job title is a flag planted for Chapter Nine) as *computers*; she was his calculator, checker and collaborator, and by every account the sharper proofreader of the two. The procedure: take a passage the guesser has never seen — the family copy of Raymond Chandler's *Pickup on Noon Street* was pulled off the shelf — and reveal it one letter at a time; before each reveal, Betty guesses the next letter; count the guesses. (The paper's published trials ran, more decorously, on a biography of Jefferson; the detective novel supplied the demonstration.) Where the text ran along its predictable grooves she was right again and again, whole stretches surrendering without a fight — because given "a small oblong reading lamp on the d—", what else is coming but *esk*? And wherever she was right, note the engineering fact: the letter on the page told her nothing she didn't already have. Its information content, *for her*, was zero. Only where she guessed wrong — at the pivots, the genuine forks, the entrances of the new — did the text actually deliver information. Averaged out, printed English was carrying only about one bit per letter against the four-and-three-quarters of the random alphabet: roughly seventy-five per cent scaffolding, this time, because Betty's model of the

language was so much richer than a table of triplets. And notice — Shannon certainly noticed — those numbers are a *measurement of the model*. Your error rate on next letters is, precisely and quantitatively, the gauge of how much of the language's structure — its grammar, its clichés, its worldly habits — you have internalized. Prediction isn't a party trick adjacent to understanding. Prediction is the *examination* understanding sits.

Shannon had already run the argument the other way round in the 1948 paper, in a passage which I regard as the most quietly momentous few inches of print in the twentieth century, though it presents itself as a bit of fun. If a statistical model can *judge* the next letter, it can *choose* one: sample from the model's own odds, then again, then again, and the model writes. He did it by hand — no computer; tables of random numbers, and for the higher orders an ingenious trick of opening a book at random and reading forward from wherever the last-chosen letters happened to occur, which makes the book itself the model, a detail I shall let you savour privately. With letters chosen at uniform random, you get static: XFOML RXKHR-JFFJUJ. Weight the letters by English frequency and a faint anglophone rhythm appears. Condition on pairs, then triplets, and pronounceable pseudo-words crawl ashore: IN NO IST LAT WHEY CRATICT FROURE. Move to whole words, chosen by their likelihood of following the previous word, and — I quote the machine of 1948, which was a man with a pencil — you get: THE HEAD AND IN FRONTAL ATTACK ON AN ENGLISH WRITER THAT THE CHARACTER OF THIS POINT IS THEREFORE ANOTHER METHOD FOR THE LETTERS THAT THE TIME OF WHO EVER TOLD THE PROBLEM FOR AN UNEXPECTED. Read that twice. It is nonsense — but it is *English* nonsense, unmistakably; it has the fist. Every step up in the model's memory bought a step up in the plausibility of its output, and Shannon says so, dryly, and moves on, leaving history's first machine-generated sentence sitting in the Bell System Technical Journal like an unexploded shell. The recipe was published. Predict better, and you

will write better. It waited fifty years for the hardware — for the explosion-drawing chips of Chapter Five and the scanned libraries of Chapter Eight — but the recipe never changed. (A few saw it early: Ray Solomonoff, one of the handful at the 1956 Dartmouth gathering where "artificial intelligence" got its name, spent his life arguing that this was the whole game — that compression and prediction and learning are one act under three names. He was politely admired and comprehensively ignored, a fate this book should by now have taught you to read as a promising sign.)

I owe you, before the punchline, the man himself, because Shannon is the rare figure in this story whose character is as instructive as his theorems. Having invented the intellectual foundation of the age by thirty-two, he declined to spend the rest of his life being its pope. He rode a unicycle through the long corridors of Bell Labs, sometimes juggling as he went; he built ever-worse unicycles, including one with an off-centre hub that bobbed the rider up and down like a duck, on the principle that anything worth doing is worth doing wrongly in an interesting way. He built THROBAC, a calculator that worked in Roman numerals; and a mechanical mouse named Theseus that could learn its way through a maze — the mouse was a magnet on wheels, the learning lived in seventy-five telephone relays under the floor, and when Theseus, dropped into a rearranged maze, unlearned the old route and found the new one, early-fifties television audiences were shown, without the term being used, a machine updating its model of the world. And on his desk he kept the Ultimate Machine, built from an idea of the young Marvin Minsky's: a small wooden casket with a single switch; flip the switch, the lid rises, a mechanical hand emerges, turns the switch off, and withdraws. That is all it does. Arthur C. Clarke saw one and wrote that there is something unspeakably sinister about a machine that does nothing — absolutely nothing — except switch itself off. Sinister, perhaps. But consider it as autobiography: the inventor of the information age kept on his desk a device whose entire repertoire is the refusal of its own function — one bit, spent perfectly. He also wrote, in

1950, a paper on programming a computer to play chess, which belongs to Chapter Five and is mentioned here only so that some reviewer cannot say I forgot it. The end was cruel in the way the gods reserve for the apt: Alzheimer's disease, which took, one bit at a time, the century's greatest theorist of memory and message. He died in 2001, just as the world his mathematics made possible was coming fully online, and by then it had rather stopped recognizing him too. He would, everyone agrees, have been delighted by the toys.



So now assemble the chapter, because the pieces click together like relays.

The Chappe telegraph taught us that a message is cheaper the more the two ends already share: send pointers, not prose. Morse and Vail, via the printer's type case, taught us that language has a countable shape, and that a code fitted to the shape beats a code that ignores it. The operators' fists taught us that even the unchosen residue of a transmission has a signature. Kelvin's cable turned communication from craft into physics; Nyquist and Hartley made the physics pay by the bit, and threw meaning out of the engineering. Markov, counting Pushkin's vowels to spike a theologian, showed that a text's structure could be captured in a table of odds on what comes next. And Shannon fused the lot: information is resolved uncertainty; English is half to three-quarters scaffolding; the quality of a mind's model of language is exactly its skill at the guessing game; and a model that can judge the next letter can be turned round to *choose* it, at which point the model writes prose in proportion to how well it has learned to predict. Every one of those clauses was in print by 1951. The talking machine was, from that moment, a matter of scale and patience.

And scale and patience are what it got. I shall keep the technical brochure for later chapters, but the essence you can now state in one sentence, and you should, at parties: the machine that talks to you was

trained by playing Betty Shannon's game (you have already seen this machine once, from the gunner's side, as Chapter Three's largest curve-fit ever attempted; here is what the curve was fitted *to*) — here is a text, cover the ending, *guess what comes next*, be scored on your surprise, adjust, repeat — played not across one evening with one detective novel but across something approaching everything ever written, trillions of rounds, with the model's table of odds grown from Markov's four pencil-tallied cells into more adjustable dials than there are words spoken in a human lifetime. The score by which it is graded is, letter for letter, Shannon's score: minimize surprise, bit by bit. That is the whole of the signal that builds it — a finishing school of human approval comes later, and a later chapter will meet it — and in the building there is no other examiner in the room. Out of nothing but that — out of being wrong about the next word a few trillion times and getting incrementally less wrong — came grammar, then idiom, then facts, then the ways of detectives and of blondes, then the ability to explain Kelvin's galvanometer to a nine-year-old in Welsh; because each of those, it turns out, is a thing you need on board in order to shave the next fraction of a bit off your error rate. Understanding wasn't the goal of the game. Understanding was the *cost of winning it*.

Which brings us to the scandal — the one people discover with a cry of triumph roughly once a week, as if it had been hidden, when it has been the entire published basis of the field since 1948. *It's just predicting the next word*. And so it is. That is precisely, exhaustively, what it is doing, and the engineers will tell you so with a candour unusual in the history of technology, and the statement is entirely true and entirely beside the point — the two conditions, I have found, in which a statement does best on the modern internet. It is beside the point because the interesting question was never *whether* it is a guessing game; the interesting question is what you are obliged to know in order to guess *that well* — and Shannon answered it in 1951: success at the game is the measure, the only measure on offer, of

how much of the world's structure you carry. "Just predicting" is doing an extraordinary amount of work in that sentence, rather like "just vibrating air" as a description of Mozart. And before we get too superior about machines that merely guess what comes next, it is worth asking what your own brain is doing when a friend pauses mid-sentence and you feel the ending arrive in your mouth before it arrives in the room; what it is doing when you read half a mangled sentence without noticing; what it was doing a moment ago, when your eye, three words ahead of your understanding as it always is, had already wagered on how this clause would—

Quite. You see the difficulty. Fifty per cent redundant, and you and I have been exploiting it since the nursery. The guessing game is not what the machine does instead of what we do. It is the part of what we do that we finally learned to measure.

But — and here the chapter turns for home — hold on to what Shannon actually built, and what he pointedly did not. His theory is a theory of the *channel*, and it is deliberately, proudly indifferent to truth. The bit count of a lie is the same as the bit count of the truth it displaces; the theorems guarantee your message arrives intact, not that it was worth sending — Thoreau's sneer at the top of this chapter is unrefuted after a hundred and seventy years. A machine reared on the guessing game inherits that indifference in its bones: it was never once scored on being right — only on being *likely*, and likely means likely-to-have-been-written, and the training library contains everything we ever wrote, the encyclopedias and the quackery, the confessions and the propaganda, in one unlabelled heap. When such a machine declared me dead in the Introduction, fluently, with a date, it was not lying; lying was beyond its pay grade. It was playing its one game honestly: an obituary is a very *probable* sort of sentence for text about an elderly broadcaster to continue into, and probability was the only virtue anyone ever taught it. The Blanc brothers, you recall, made their fortune the same way — by arbitraging the gap between what travels fast and what happens to be true — and the law

arrived only after the profits were banked. That gap is now the most important terrain in public life, and Chapter Ten is camped on it, and I shall meet you there.

In the meantime, the guessing game had one more startling property, and it sets up everything that happens next in our story. Nothing in it — look back through the chapter; check me — nothing in it is about *English*. Markov counted Russian. The game is indifferent to the tongue it is played in, as indifferent as the towers were to the dispatches, and a machine that can learn the odds of one language can learn the odds of another, and if it can learn them both at once, in the same table, then something very strange becomes possible in the space between them. Which is why the next chapter begins with an army of Napoleon's engineers, digging a trench in the Egyptian delta in 1799, and striking, with a clank that echoed for two centuries, a broken slab of black stone with the same message written on it three times.

Getting the message, it turns out, was the easy part. Now we have to work out what it says.

CHAPTER SEVEN

Lost in Translation

When I look at an article in Russian, I say: "This is really written in English, but it has been coded in some strange symbols. I will now proceed to decode."

— WARREN WEAVER, LETTER TO NORBERT WIENER,
1947

In the middle of July 1799, at a ruined fort on the westernmost mouth of the Nile near the town the locals called Rashid and the Europeans called Rosetta, a party of French soldiers was doing what soldiers spend most of their time doing, which is digging. Napoleon had invaded Egypt the summer before, with thirty-eight thousand troops and — because this was Napoleon — a hundred and sixty-seven savants: mathematicians, naturalists, chemists, artists and engineers, brought along to measure, catalogue and generally annex the intellectual property of the pharaohs while the army annexed everything else. The expedition had not gone

entirely to plan. Nelson had found the French fleet at anchor in Aboukir Bay and removed it from the board in a single evening, which left the army of the Orient stranded in the Orient, and so the French dug in, literally, strengthening every strongpoint they held. At Rosetta that meant demolishing the crumbling walls of an old Ottoman fort to build a better one, and it was in the rubble of that demolition that an officer of engineers named Pierre-François Bouchard noticed a broken slab of dark, dense stone, about the size of a coffee table, covered edge to edge in writing. Three bands of writing, in fact, in three different scripts. And Bouchard, who like every product of the French engineering schools had been made to suffer enough Greek to recognize it when he saw it, looked at the bottom band and realized that whatever the top of the stone said in the mysterious picture-writing of ancient Egypt, the bottom of the stone said in a language educated Europe could read.

The stone went to Cairo, to the institute the savants had set up in a commandeered palace, and the savants understood at once what they had, because the Greek text helpfully explains it. The slab is the bottom third of a temple stele recording a decree issued at Memphis in 196 BC by a general council of Egyptian priests, in honour of the thirteen-year-old Greek pharaoh Ptolemy V, on the anniversary of his coronation. The content, I'm sorry to report, is not the wisdom of the ages. It is a press release. The boy king has cut the temples' taxes, forgiven certain debts, released certain prisoners, restored certain privileges to the priesthood, and in return the priests decree that his statue shall stand in every temple and his festival shall be celebrated with the appropriate enthusiasm. Bureaucracy, flattery and tax relief: three thousand years of civilization and the paperwork is instantly recognizable. But the genuinely important part of the Rosetta Stone — and I want you to hold on to this, because the whole chapter is hiding inside it — is the small print at the end. The decree closes by ordering that it shall be inscribed on stelae of hard stone *in sacred script, in native script, and in Greek letters*, and set up in every major temple in

Egypt. The stone is an instruction to copy the stone. One meaning, three encodings, replicated across a country. The priests of Memphis, wanting only to flatter a teenager, had manufactured the one artefact civilization would need, two thousand years later, to get its memory back: a parallel text.

Why the memory needed getting back is a story of its own. Hieroglyphic writing did not die with the pharaohs; it died with the priests, slowly, under Rome. The last datable hieroglyphic inscription was carved at the temple of Philae in AD 394, and when the last priest who could read the old script died, some time not long after, the entire written record of the oldest civilization on Earth — the temple walls, the obelisks, the acres of papyrus — fell silent, all at once, like a library at closing time. Worse than silent: misread. A late-antique treatise called the *Hieroglyphica*, attributed to one Horapollon, turned up in 1419, on a Greek island, in the luggage of a Florentine traveller, and assured European scholars that hieroglyphs were not writing at all but pure mystical symbols — the goose means "son", the hare means "open", each sign a little sealed capsule of philosophy — and Europe, which adores a mystery, believed it for four hundred years. The magnificent Jesuit polymath Athanasius Kircher, a man wrong about more things more confidently than almost anyone in the seventeenth century, published whole volumes translating obelisks into paragraphs of Neoplatonic profundity — signs which, we now know, mostly just spell out the name and titles of a king. (In fairness to Kircher, he also championed the study of Coptic, the language of Egypt's Christians, and insisted it was descended from the language of the pharaohs. He was right about that, and the fact turns out to matter enormously, so let us allow the old fantasist his one bull's-eye.)

Then came the diggers of Rosetta, and then, promptly, the British. In 1801 the stranded French army surrendered, and under Article 16 of the Capitulation of Alexandria the antiquities the savants had collected passed to the victors. The French commander, General Menou, tried to keep the

stone by declaring it his personal property, an argument the British received with the sympathy it deserved; the savants for their part threatened to burn their notes and specimens rather than hand them over, remarking that history would remember who had made the second bonfire of Alexandria — and the British, to their credit, let them keep the papers and took the rock. It was hauled aboard a frigate called HMS *Egyptienne* — a captured French ship, because the age had a sense of humour — and it has sat in the British Museum since 1802, where for many years its side carried the painted inscription CAPTURED IN EGYPT BY THE BRITISH ARMY IN 1801. It is a war trophy, and one may as well say so plainly; though say what you like about the British, they do at least label the loot.

But here is the detail that matters, the one the treaty-writers missed. Before the surrender, the French had inked the stone's face and pressed paper onto it, using the most famous object in the history of writing as its own printing plate, and the copies were already in Paris. The British had possession of the granodiorite. Nobody had possession of the *text*, because a text is not a thing you can possess by holding it; the information had detached itself from the object and gone travelling, as information does. Which meant that the race to decipher it would be run on paper, in London and Paris simultaneously, by two nations recently at war and permanently at odds — and if you wanted to design a competition to crack the code of the pharaohs, you could hardly do better than a Frenchman and an Englishman who couldn't stand losing to each other.



The Englishman was Thomas Young, and if the name rings only a faint bell, that is because he committed the career error of being brilliant at everything, so that every field he transformed remembers him only as a visitor. Physician by trade; the man whose double-slit experiment showed that light behaves as a wave; the Young of Young's modulus, which is why engineers know him; proposer of the three-colour theory of vision, which

is why your television works; and, in an idle moment in 1813, coiner of the term "Indo-European", which is why an entire science of language history knows him. A later biographer called him the last man who knew everything, and the striking thing is how little anyone has disputed it. Young came at the stone the way he came at everything, as a puzzle in pattern. He counted. The name Ptolemy appears in the Greek text a certain number of times; among the hieroglyphs, certain groups of signs recur, ringed by an oval border the French called a *cartouche*, because it looked like a rifle cartridge; match the frequencies, line up the recurrences, and you can begin to guess that the cartouche spells the king's name — and if it spells it, then at least some hieroglyphs are not mystic capsules but humble letters, standing for sounds. By 1819 Young had extracted a handful of genuine sound-values and published them, anonymously, in a supplement to the *Encyclopaedia Britannica*. And then he stopped, halted by the very reasonableness that made him. Foreign names like Ptolemy, he reasoned, had to be spelled out phonetically, since Egyptian picture-writing could hardly contain a native sign meaning "Macedonian Greek"; but the script as a whole, he remained sure, was symbolic. He had the key in the lock and declined, on principle, to turn it.

The Frenchman turned it. Jean-François Champollion was everything Young was not: poor, provincial, monomaniac, a man who knew one thing — but it was the right thing, and he had known it since childhood. Born in Figeac in 1790, he was a linguistic prodigy of the kind that makes teachers nervous, and at the age of about eleven he was shown a collection of Egyptian antiquities by the prefect of his district, who happened to be Joseph Fourier, the mathematician who had gone on the Egyptian expedition and whose analysis of heat flow, by one of those long sideways throws this book keeps making, underlies half the signal-processing on the planet, including the compression of the recording of this chapter, should you be listening to one. The boy stared at the unreadable signs and, the story goes, resolved on the spot that he would be the one to read them. The

story is probably burnished — origin stories always are — but what is documented is what he did next, which was to spend his adolescence acquiring Hebrew, Arabic, Syriac, Chaldean, Persian and above all Coptic, the language Kircher had correctly fingered as ancient Egyptian's living descendant, still chanted in the liturgy of Egypt's Christians. Champollion learned Coptic the way you learn a language you are in love with; he kept notebooks in it and claimed to talk to himself in it, which must have enlivened the streets of Grenoble. His logic was simple and turned out to be everything: if the old script wrote the old language, and the old language survived in the new, then a man who could sound out the signs might *hear Egyptian again* — and check his guesses against a tongue still spoken.

The break came in two moves. First, a second parallel text: an obelisk shipped home to Dorset by the English collector William Bankes, bearing a Greek inscription with the names Ptolemy and Cleopatra, and cartouches to match. The two names share letters — P, O, L, T — and if the script were phonetic, the shared sounds should appear as shared signs, in the right positions, in both cartouches. They did. That confirmed the method for Greek names. Then, on the fourteenth of September 1822, fresh drawings arrived in Paris from the temples of Abu Simbel — deep Egypt, old Egypt, carved a thousand years before any Ptolemy — and Champollion, examining a royal cartouche, saw a sun disc, a sign he could not yet read, and a doubled sign he suspected was S. The Coptic for sun is *re*. Re — something — s — s. *Ramses*. A second cartouche gave him an ibis, the bird of the god Thoth, then the same endings: *Thutmose*. The names of the greatest pharaohs, spelled by sound, in native Egyptian, using pictures as letters and letters as pictures — the script was not symbolic *or* phonetic, it was both at once, a rebus the size of a civilization. Champollion ran from his rooms to the Institut, where his elder brother worked, burst in on him, cried "*Je tiens l'affaire!*" — I've got it! — and fainted. By the family's account he then took to his bed for the better part of five days, poleaxed by his own discovery, and I pass this on with delight, because in a book full of

overclaiming it is restful to meet the one man in the history of the subject whose reaction to a breakthrough was physical collapse rather than a press conference. Two weeks later, on the twenty-seventh of September 1822, his *Lettre à M. Dacier* was read before the Académie des Inscriptions in Paris, laying out the phonetic system. In the audience, on a visit from London, sat Thomas Young.

You can guess the aftermath. Young, with genuine grace and genuine grievance mixed in the English manner, congratulated the discovery and spent his remaining years suggesting it had been half his; French and British partisans have refought the question ever since, with the score standing roughly at: Young found the door, Champollion walked through it, and the walking is the discovery. Britain, meanwhile, kept the stone and lost the race — the object in London, the meaning in Paris, which tells you precisely how much of the meaning lives in the object. Champollion got his chair, the first professorship of Egyptology in the world, created for him at the Collège de France, enjoyed it for barely a year, and died of a stroke in 1832 at the age of forty-one, worn through, his grammar of Egyptian still in manuscript. His brother published it. The dead language has been talking ever since.

Now stand back from the anecdote, as this book keeps asking you to do, and look at what actually happened, because it is one of the most remarkable proofs of concept in intellectual history. A civilization's entire semantic content was recovered — not guessed at, *recovered*, checkably — by two men who had never heard a word of the language spoken, from pattern, frequency and a few square feet of parallel text, plus one lucky survival in a church liturgy. No teacher. No phrasebook. No grammar, until Champollion wrote it himself, afterwards, as a *description of what he had already decoded*. The lesson, which the nineteenth century filed under romance and the twentieth would have to relearn at industrial expense, is this: meaning survives re-encoding, and given enough aligned text and enough patience, it can be pulled back out of the pattern alone. Align,

count, guess, check. Hold on to that method. It returns at the end of this chapter, wearing silicon, and it will have read everything.



Of course, the truly ambitious have never wanted to *translate* Babel. They have wanted to repeal it.

The book of Genesis, you'll recall, is admirably clear about the mechanics: the whole earth was of one language and one speech, mankind used the resulting perfect coordination to start building a tower with its top in the heavens, and the Almighty, taking the point at once — "now nothing will be restrained from them, which they have imagined to do" — confounded their language, that they may not understand one another's speech. Linguistic diversity as a security measure. What matters here is the itch the story left behind: for two thousand years, some of the best minds in Europe were convinced that beneath the rubble of Babel the original language was still down there somewhere, and that finding it — or failing that, rebuilding it — would restore mankind's lost operating manual. The recovery attempts ranged from the scholarly to the frankly experimental. Herodotus tells of the pharaoh Psammetichus, who had two infants raised in silence, to discover what language would emerge unprompted; their first recognizable word was, allegedly, *bekos* — Phrygian for bread — whereupon Egypt, with commendable empiricism, conceded linguistic seniority to Phrygia. James IV of Scotland is said to have repeated the experiment on an island in the Firth of Forth, and the children, a chronicler assures us, grew up speaking very good Hebrew. The chronicler was not present. Neither, one suspects, was the Hebrew.

The rebuilding attempts were more serious, and one of them was very serious indeed. In 1668 John Wilkins — founding spirit of the Royal Society, sometime bishop of Chester, and a man who had survived being married to Oliver Cromwell's sister, so no stranger to ambitious projects — published *An Essay Towards a Real Character, and a Philosophical Lan-*

guage, six hundred folio pages proposing to replace all existing languages with a rational one. The scheme is beautiful in the way that only doomed schemes are: divide everything in the universe into forty fundamental categories, subdivide down to species, and let every word *be* its own address in the tree of knowledge, so that the name of a thing tells you what the thing is. No synonyms, no ambiguity, no Babel. The difficulty, as commentators from the Royal Society down to Jorge Luis Borges have enjoyed pointing out, is that the universe declines to hold still in forty boxes; every taxonomy is a portrait of its taxonomist, and a language that can only say what the classification permits is a prison pretending to be a filing system. Leibniz, a generation later, dreamed one floor higher: his *characteristica universalis* would render thoughts themselves in exact symbols, so that reasoning became calculation, and — his own delicious phrase — when two philosophers disagreed, they would need no more debate than two accountants: they would sit down with their pencils and say to one another, *let us calculate*. Hold on to Leibniz: a machine that turns language into arithmetic is precisely what this chapter is on its way to building — except that it will take the messiness on board instead of legislating it away.

The last great assault on Babel was launched from inside it. Ludwik Lejzer Zamenhof grew up in the 1860s in Białystok, a city of the Russian Empire where Poles, Russians, Germans and Jews lived side by side in four languages and mutual detestation, and the young Zamenhof — Jewish, therefore fluent in the detestation from the receiving end — concluded that the languages were the cause and a common one would be the cure. He became an eye doctor, which feels almost too neat, and in 1887 published, under the pseudonym *Doktoro Esperanto*, "the doctor who hopes", a slim brochure containing a complete invented language: sixteen grammatical rules, no exceptions, vocabulary quarried out of the common roots of Europe. Esperanto is the most successful constructed language in history — banned by Hitler and Stalin both, which is a distinction of a kind — and it

settles the argument all the same, because after a century and a half the number of human beings who speak it remains a rounding error on the number born every year speaking something else. The dream fails, every time, for the same reason: a language is not a nomenclature bolted onto the world, it is a living record of everything its speakers have found worth distinguishing, and no committee can issue that. You cannot repeal Babel. The towers stand.

Which leaves the other option, the humble one: leave the towers standing and dig tunnels between them, one pair at a time — translation, the second-oldest profession of the word trade. For most of history it was slow, expensive, artisanal work, and there matters rested, until the twentieth century did what the twentieth century did to everything artisanal. It is one of the reliable ironies of this book that the money to mechanize understanding between peoples came from the business of destroying them, and translation is no exception. The tunnel-diggers' great subsidy arrived in uniform.



Cryptanalysis is translation's dark twin — the same task with an uncooperative author. A ciphered message is simply a text in a language designed to have no second speaker; the codebreaker is a translator working against an adversary instead of a dictionary. And in the Second World War, at a Victorian estate north of London called Bletchley Park, the twin got, for the first time in history, an industrial budget. You met Alan Turing two chapters ago at the chessboard, wondering whether a machine could imitate your aunt; here he is in his other capacity, hunched over the traffic of the German Enigma machine, and I shall be brief, because his story is told elsewhere and only one feature of Bletchley belongs in this chapter. It is this: the great engine of the place, the electromechanical Bombe that Turing designed and Gordon Welchman improved, did not read German and did not break codes by insight. It worked by *guess*

and check — by taking a crib, a scrap of guessed plaintext, and testing rotor settings by the thousand until the guess and the cipher could be made to line up. And where did the cribs come from? From German habit. Weather stations transmitted, every morning, in rigid formulaic phrasing, reports containing the word WETTER; quiet sectors dutifully reported KEINE BESONDEREN EREIGNISSE, nothing to report, day after obliging day. A crib is a small, stolen, single-use Rosetta Stone: a parallel text, one line long. Bletchley was Champollion run as heavy industry — alignment, frequency, guess and check, executed by machinery at a speed no human could match, on a deadline measured in convoys. It proved, behind the highest wall of secrecy in the world, that language yields to statistics plus machinery plus parallel text. The war ended, the machinery was smashed on orders, the secret held for thirty years — and the idea walked out anyway, in the heads of the people who had watched it work.

One of the people watching from the American side of that secret world was a mathematician and foundation officer named Warren Weaver — the man who packaged Shannon's mathematical theory of communication for the general reader; during the war he had run the Applied Mathematics Panel, brokering the mathematics of fire control and statistics to whoever needed it, and he had seen what the cryptanalysts could do. In March 1947 he wrote to Norbert Wiener, the founder of cybernetics, with a modest proposal, and one sentence of that letter has become the founding text of an entire field — you will find it at the head of this chapter. When I look at an article in Russian, I say: this is really written in English, but it has been coded in some strange symbols; I will now proceed to decode. Translation as cryptanalysis. The foreign is just the familiar, encrypted. Wiener, to his lasting credit as a linguist and lasting embarrassment as a prophet, poured cold water on the whole idea — the boundaries of words were too vague, he replied, the emotional connotations too extensive, for any quasi-mechanical scheme to be hopeful. Weaver waited two years and then went over everyone's head in the politest

way possible: in July 1949 he wrote a memorandum, titled simply "Translation", and posted it to some two hundred of the best-connected scientists in America. Money listens when Weaver talks — he directed the Rockefeller Foundation's science division, and had an unmatched nose for the about-to-be-possible — and the memorandum is a genuinely strange and wonderful document, half engineering prospectus, half prose poem. It proposed logical analysis, cryptographic methods, statistical regularities. And it closed with a metaphor worth the whole of the rest: think of individuals living in a series of tall closed towers, all erected over a common foundation, shouting to each other from tower to tower in vain. But descend, said Weaver, and the towers open onto a great open basement, common to all. Might not the way to translate be not to shout between the towers but to go down — beneath Chinese, beneath Russian, beneath English — into the deep shared basement of human communication, and walk across?

Remember the basement. It will take seventy years, and it will not look remotely like what Weaver imagined, but the field he founded with a mailing list is going to end up down there.



What turned a memorandum into an industry was, as usual, fear. In August 1949 the Soviet Union tested its first atomic bomb, years ahead of Western expectation, and Washington woke up to an uncomfortable arithmetic: the enemy published. Soviet science poured out journals — physics, chemistry, metallurgy, aeronautics — a mountain of open literature growing faster than the tiny corps of scientifically literate Russian translators could ever climb. Somewhere in that mountain, presumably, were the clues to the next unpleasant surprise. The choice was stark: train ten thousand translators, or build a machine. Congressmen understood machines. The money began to move — Air Force money, Navy money, CIA money laundered through respectable foundations —

and by 1952 the new field could hold its first conference, at MIT, convened by a philosopher named Yehoshua Bar-Hillel who holds the curious distinction of being the first human being ever employed full-time on machine translation. File his name; he returns shortly with a hatchet.

Then came the demonstration, and the demonstration is the part of the story everyone should be made to study, because it set the pattern for every artificial-intelligence announcement since. On the seventh of January 1954, at IBM headquarters on Madison Avenue, a collaboration between IBM and Georgetown University put on a show for the press. An IBM 701 — the company's first commercial scientific computer, a machine ordinarily employed on missile trajectories — had been loaded with a Russian vocabulary of exactly two hundred and fifty words and precisely six rules of grammar. An operator who spoke no Russian punched romanized Russi-

an sentences onto cards; the machine clattered; and out came English. *Mi pyeryedayem mislyi posryedstvom ryechyi*: "We transmit thoughts by means of speech." The quality of coal is determined by calorie content. Starch is produced by mechanical methods from potatoes. Sixty-odd sentences, hand-picked from chemistry and general affairs, every one rehearsed, the whole repertoire chosen with loving care to stay inside the two hundred and fifty words and the six rules — and it worked, and the press went up like a flare. Electronic brain translates Russian; the language barrier falls. Léon Dostert, the Georgetown linguist running the show — a man who had been Eisenhower's wartime interpreter and had organized the simultaneous-translation system at Nuremberg, so nobody's fool about how hard translation actually is — announced that meaningful machine conversion of whole languages might well be an accomplished fact within five years, perhaps three. Everyone involved knew the demonstration was a conjuring trick, in the honest theatrical sense: nothing up the sleeve except the fact that the deck contained only sixty cards. Nobody said so into a microphone. The funding agencies read the headlines, the Soviets read the

headlines and started their own programme, which produced fresh American headlines about a translation gap, and then in 1957 Sputnik went over everyone's heads at eighteen thousand miles an hour, beeping in Russian, and the money question was settled for a decade. By the early sixties, machine translation was consuming millions a year across a dozen laboratories, and the five years were always five years away.

It is at this point in the story — the point of maximum promise — that folklore inserts the most famous anecdote in the history of the field, and I would not dream of denying it to you. The engineers, the story goes, fed their machine the English sentence "The spirit is willing but the flesh is weak," had it translated into Russian, then fed the Russian back for translation into English, and the machine returned: "The vodka is good but the meat is rotten." A sister version does the same to "out of sight, out of mind" and receives "invisible idiot". Now: it never happened. Almost certainly. No one has ever produced the machine, the date, the laboratory or the printout; the tale was already circulating in the fifties attached to systems that could not yet translate anything into anything, let alone back; and the historians of the field who went looking for the source found only earlier versions of the same joke, wearing earlier hats — it appears to be descended from a family of human-interpreter jokes that had been doing the rounds since before the computer existed. It is, in short, apocryphal, and I flag it with relish rather than regret, because the vodka story is that most useful kind of untruth: a fable the field told about itself, which turned out to be more accurate than its press releases. The machine of 1960 could not manage idiom, metaphor, scripture or context, and everyone in the laboratories knew it, and the joke was how the knowledge got past the funding agencies and out into the open. The anecdote is false; the moral is certified. It remains the perfect epigraph for the entire genre of the overpromising demo, which is why I have hung it, in spirit, over this chapter — the spirit being willing, and the flesh, as we shall now see, being weak.

The first blow came from inside. Bar-Hillel — the first man hired into the field, remember — published in 1960 a survey of the whole enterprise which is one of the great acts of intellectual honesty of the century, and its centrepiece is a sentence of childlike simplicity: *The box was in the pen*. Every English speaker reads it instantly: the pen must be a playpen, since boxes don't fit inside fountain pens. But nothing in the dictionary tells you that, and no rule of grammar tells you that; what tells you is your knowledge of boxes, pens, children and the sizes of things — your model of the world. Fully automatic high-quality translation, Bar-Hillel concluded, was not around the corner; it was not even on the map, because translation is not a linguistic task with occasional excursions into knowledge — it is a knowledge task wearing linguistic clothing. The field's founder had discovered, six years into the five-year promise, that the problem contained the whole world. He recommended aiming lower. Nobody likes being told to aim lower.

The second blow came from the auditors. In 1964 the funding agencies, noticing that the five years had now been five years away for ten years, convened the Automatic Language Processing Advisory Committee — ALPAC, an acronym that still makes elderly computational linguists flinch — under the chairmanship of John R. Pierce of Bell Labs, the engineer who had given the transistor its name. Pierce was the ideal chairman from every point of view except that of the people being audited: a first-rate engineer, a friend of ambitious technology, and constitutionally allergic to glamour masquerading as progress. The committee's 1966 report is short, dry and lethal. Machine translation, it found, was slower, less accurate and roughly twice as expensive as human translation; there was no shortage of human translators for the government's actual needs; some twenty million dollars had been spent; and there was — the sentence that closed a hundred laboratory doors — no immediate or predictable prospect of useful machine translation. The report recommended funding computational linguistics as *science*, modestly, and machine translation as

a *product* not at all. Washington obliged with enthusiasm. Within two years the field was a ghost town; researchers scattered into other disciplines or other countries; and the whole cycle — dazzling demo, ecstatic headlines, five-year promise, quiet audit, sudden freeze — completed its first full revolution. The pattern would repeat, and has repeated since, often enough that it now has a name: the AI winter. It is worth recording, for the connection-minded, that the first one arrived by way of translation. The field that had founded machine intelligence's public reputation was also the first to bankrupt it — and the deep cause, both times, was the same: everyone had mistaken language for a code, when language, as Bar-Hillel's little box-in-the-pen had demonstrated, is a code whose key is everything you know.

Not quite everything died in the frost. The Air Force still needed its Soviet journals read, whatever the auditors said, and a Hungarian-born engineer named Peter Toma kept a rule-based system alive that became SYSTRAN, which spent the Cold War quietly grinding Soviet technical literature into rough but usable English for the American military — the founding purpose of the whole field, surviving in a basement while the theory wept upstairs. SYSTRAN translated for the Apollo-Soyuz mission, was adopted by the European Commission in 1976, and thirty years after ALPAC was still on duty as the engine behind the web's first free translator, which its owners called Babel Fish, after the creature in Douglas Adams that you put in your ear — the dream of Babel reduced, at last, to a brand name. And in Montreal, the following year, a university group switched on a system called Météo that translated Canadian weather bulletins between English and French — millions of words a year, decade after decade, so reliably that human translators were relieved to surrender the work, weather bulletins being the kind of prose that wears out the soul. Météo worked because a weather bulletin is a language with the ambition surgically removed: a few hundred words, a handful of sentence shapes, no metaphors, no vodka. That was the lesson the survivors drew: cage the

language and the rules will hold it. The lesson underneath, which took another decade to sink in, was grimmer. Rules do not scale. Every rule breeds exceptions, every exception breeds sub-rules, and real language — the kind with jokes, ellipses, sarcasm and this morning's brand-new slang in it — is not a tidy system with some untidy corners. It is exceptions all the way down, with a rule problem.

★

The thaw, when it came, came from people who were not translators at all, did not especially care about linguistics, and had spent their careers listening to noise. And it began, as the best parts of this book do, with a man history had thrown across several languages the hard way. Frederick Jelinek was born in Prague in 1932; his father, a Jewish dentist, died in the Theresienstadt concentration camp; the family got out in 1949, and the boy arrived in New York with no English, acquired it, and set about becoming an electrical engineer, having judged that a refugee's accent would be less of a handicap in engineering than in the law career he wanted. By the 1970s and 80s he was running the continuous-speech-recognition group at IBM's Watson Research Center, and the group's heresy was total. The reigning approach to speech, as to translation, was to teach the machine linguistics: phonetics, syntax, rules. Jelinek's group threw the rules overboard and treated speech as Shannon had treated everything — as a signal in a noisy channel, to be recovered by probability. What is the most likely word, given the sound and given the words so far? Don't ask a linguist; count. Feed the machine text by the millions of words and let the frequencies speak. It worked — it worked better every time the corpus got bigger — and out of that group came the most quoted sentence in the history of the field: *Every time I fire a linguist, the performance of the speech recognizer goes up*. I am obliged to report, in the now-familiar ritual of this chapter, that Jelinek variously denied saying it, or said he could not remember saying it, or observed that what he actually said was

gentler; and I am equally obliged to report that the sentence has outlived every correction, because — like the vodka — it is apocrypha doing honest work. Something true had been noticed, and it was scandalous: at sufficient scale, counting beats understanding. Or rather — hold the thought, because the whole book turns on it — at sufficient scale, counting *becomes indistinguishable from* understanding, and no one has yet located the join.

Around 1988, Jelinek's statisticians — Peter Brown, Robert Mercer, the brothers Stephen and Vincent Della Pietra, names to remember for the strangest of reasons in a moment — wandered from speech into translation, carrying the noisy-channel idea with them, and what they did was resurrect Warren Weaver's forty-year-old joke as an equation. A French sentence, they proposed with straight faces, is merely an English sentence that has been garbled in transmission through a noisy channel; translation is the recovery of the original. This is, of course, preposterous as a theory of French — the French have been claiming precisely the reverse for centuries — but as an engineering stance it changes everything, because it splits translation into two problems statistics knows how to attack: which English sentences are probable at all (count English), and which English words tend to line up with which French words (count *pairs*). And there was the rub: to count pairs, you need parallel text — not a broken stele's worth but mountains of it, sentence after sentence in faithful alignment, more than any team could ever commission. They needed to find it lying about, already made, as a by-product of something else. They found it in Canada. By law, everything said in the Canadian Parliament is published in both English and French: Hansard, the verbatim record, translated overnight, year upon year upon decade — millions of sentence pairs of impeccable alignment and magnificent tedium, the collected complete works of Canadian democracy asking itself supplementary questions about fisheries. The IBM team got the tapes. Let the reader savour the shape of it, because it is this book's shape in miniature: the Rosetta Stone that revived machine translation was manufactured by parliamentary procedure —

bureaucracy's exhaust, warehoused for constitutional reasons, waiting like Jacquard's punched cards and the video-gamers' explosion chips for somebody to notice it was treasure. Nobody builds the future's raw material on purpose. It accumulates, in triplicate, in the files.

The system they built was called *Candide* — somebody in Yorktown Heights had a sense of humour about best-of-all-possible-worlds optimism — and it learned to translate in a way that would have made every linguist since Pāṇini reach for the smelling salts: from nothing but the aligned sentences. No dictionary, no grammar, no rules of French. The machine discovered for itself, by iterating over the millions of pairs, that *not* keeps company with *ne... pas*, that adjectives swap sides of their nouns on the way across the Channel, that one word may become three and three one. Where Champollion aligned a few hundred cartouches by hand and by genius, *Candide* aligned millions by arithmetic and let the probabilities do the noticing. The linguistic establishment's reaction may be politely described as horror. Statistics, the Chomskyan mainstream had long held, was beneath the dignity of the subject — Chomsky himself had written, years before, that the notion "probability of a sentence" is an entirely useless one, under any known interpretation of this term — and here were engineers with no theory of language at all, publishing translations. Reviewers informed the IBM team that brute computation was not science. The evaluations, unmoved, kept coming back the same way: the counting worked, and the counting kept improving as the corpora grew, and within a decade the heresy was the curriculum. Statistical machine translation conquered the field so completely that a graduate student of 2005 could go a whole career without writing a grammar rule.

And the statisticians themselves? Here the story exits sideways through a door nobody saw. In 1993, Peter Brown and Robert Mercer left IBM together — poached by a mathematician named James Simons for a secretive Long Island hedge fund called Renaissance Technologies, on the theory that people who could find the signal in the babble of two languages

might find it in the babble of the markets. The theory was correct. The fund's Medallion vehicle became the most profitable money machine in the history of finance, Brown and Mercer eventually ran the firm, and the mathematics of translation quietly spent the next quarter-century translating the market's past into its future at forty-odd per cent a year. Mercer's later career as a political financier had consequences for Anglo-American public life that belong in somebody else's book, and possibly in somebody else's courtroom; I note merely, for the pattern's sake, that the trail from a Rockefeller memorandum about Russian journals leads, by way of the Canadian Hansard, to both the modern translation engine and the modern quant fund, and that nobody planned either. They never do.

The web then did to statistical translation what it does to everything: fed it. In 2004 Google hired Franz Josef Och, whose systems had been winning the government translation evaluations, and gave him the two things Google had beyond any institution in history — computers, and text. The parallel corpora this time were the United Nations, whose proceedings exist in six official languages, and the European Union, whose *acquis* and parliamentary record exist in more than twenty; once again the machine went to school on the exhaust of bureaucracy, the most boring prose on Earth teaching silicon to speak in tongues. Google Translate launched in 2006, and Och's group demonstrated over and over the lesson Jelinek had learned on speech: more data beats better algorithms, and much more data beats much better algorithms. By the 2010s the thing was serving hundreds of millions of people across dozens of languages — the tourist's phrasebook, the immigrant's form-filler, the teenager's homework accomplice — a genuine wonder of the world, and everyone who used it knew exactly how far to trust it. It produced the *gist*. Word by word, phrase by phrase, it stitched across; sentence by sentence it wobbled; paragraph by paragraph it wandered off; and anything that depended on the whole utterance meaning one thing — irony, argument, poetry, the box in the pen — went to pieces, because phrase-statistics see language

through a slot a few words wide. The field had built a magnificent tunnel network between Weaver's towers, and every tunnel was four feet high. What it needed was the basement. What it needed, though nobody put it this way at the time, was a machine that could read the whole sentence before opening its mouth.



The machine arrived in 2014, and it came — the reader will by now expect this — from the neural-network believers of Chapter Four, emerging from their long winter with the explosion chips of Chapter Five humming under their desks. The recipe was called sequence-to-sequence, and its elegance was almost insolent: one network, the encoder, reads the source sentence word by word and compresses everything it has understood into a single fixed-length vector — a list of a few hundred numbers, a point in space; a second network, the decoder, takes that point and unwinds it, word by word, into the other language. Reading as inhaling, translation as exhaling, and the entire meaning of the sentence held, for one moment between breaths, as a single position in a geometry of hundreds of dimensions. Here — pause and admire it — is Weaver's basement, actually built: beneath the words, a shared space where the sentence exists as pure position, before it has decided which language to be. And the thing half worked. On short sentences it produced translations of a fluency the statistical systems never touched. On long sentences it fell apart, and everyone knew why. A fixed-size vector is a fixed-size suitcase, and a forty-word sentence with sub-clauses does not fold. A computational linguist named Ray Mooney summarized the objection with a precision I can only bowdlerize: you cannot cram the meaning of a whole %&#!\$# sentence into a single \$&#!\$# vector. *The reader will recognize the shape of the complaint. It is the bottleneck problem, the oldest problem in this chapter: the meaning is bigger than the channel. The encoder was being asked to do to War and Peace* what the six rules of Georgetown did to Russian — and*

it responded the way every overpromising translator since 1954 had responded, by doing the first clause beautifully and improvising the rest.

The fix was proposed within months, by a master's student. Dzmitry Bahdanau, working in Yoshua Bengio's laboratory in Montreal with Kyunghyun Cho, suggested in September 2014 that the decoder be released from the suitcase: instead of translating from one crushed summary, let the machine, at every single word it produces, *look back at the entire source sentence* and learn — from the data, not from rules — where to look. Producing the French for "bank"? Glance back: is there river weather anywhere in the English? Producing a pronoun? Glance back at who is being talked about. They called the mechanism *attention*, a rare instance of a technical term that means what it says, and when researchers drew pictures of what the trained network was attending to, out came maps that would have stopped Champollion's heart: *ne... pas* linked across to "not", adjectives and nouns crossing over each other in mid-Channel swap, each output word reaching back along bright lines to the source words that justified it — the machine's eye flicking between the Greek band and the hieroglyphic band, cartouche to name, exactly the motion, learned by nothing but alignment and arithmetic. Google rebuilt Translate around neural networks with attention and threw the switch in November 2016, replacing the statistical engine overnight for its biggest language pairs, and the leap in quality was, by the company's measurements, greater than everything achieved in the previous decade combined; users noticed by morning. Somewhere the ghost of ALPAC audited the figures and, one hopes, allowed itself a thin smile.

Then came the paper. In June 2017, eight researchers at Google — Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, Polosukhin; the paper carries a footnote declaring the listing order random, everyone having contributed equally, which is either the most collegial footnote in computer science or the most hard-won — published the machine-translation paper with the cheekiest title in the history of the literature:

"Attention Is All You Need". (The title is said to be a wink at the Beatles, and I choose to believe it, on the grounds that a field this fond of apocrypha owes itself one true joke.) The argument was as blunt as the title. Attention had been bolted onto the reading-and-breathing networks as a corrective; very well — throw away the networks and keep the correction. No more reading the sentence one word at a time, no more suitcase: let every word in the sentence attend directly to every other word, all pairs, all at once, in stacked layers of mutual regard, and let that *be* the machine. They called the architecture the Transformer, and it carried the day on the standard benchmarks — English into German, English into French, the same pairs the field had been drilling since the Hansard days — at a fraction of the training cost of its rivals. And buried in that word "fraction" was the property that changed the world, though the paper spends barely a sentence preening about it. Because every word attends to every word *simultaneously*, nothing waits for anything: the whole calculation is the kind of embarrassingly parallel arithmetic that the video-gamers' explosion chip had spent thirty years being accidentally perfected for. The old networks read like a human, in order, and so trained at the speed of sequence; the Transformer reads like a page lit all at once, and trains at the speed of silicon. The flagship model of the paper trained in three and a half days on eight GPUs. The architecture was not merely better. It was *scalable* — you could make it bigger, and feed it more, and it just kept getting better, no ceiling yet visible — and if the reader has retained one lesson from Jelinek and Och and the whole counting heresy, it is what happens next to an architecture like that. You will have noticed what the T stands for in GPT.

For the paper's authors, translation was the task; for history, translation was the excuse. Within eighteen months the encoder half of the Transformer had been carved off to become BERT, reading the web for Google's search engine; the decoder half had been carved off by a laboratory called OpenAI to become the GPT series, and the decoder half is the punchline of this book, because a decoder is a machine that translates a

text's past into its future — that takes everything said so far and produces the most probable next word, which is to say Shannon's guessing game from the last chapter, industrialized — you have already had this reveal once, in Betty Shannon's parlour; here it is again, wearing the machinery. And at that point the fifty-year-old boundary of the field quietly dissolved, because it turned out that *everything is translation*. Question into answer. Document into summary. English into Python. Description into diagram; prompt into essay; muddle into plan. The tunnel-boring machine built for the single job of getting from French to English had broken through into Weaver's basement, the great open floor under all the towers — and down there, it transpired, the corridors connect everything to everything, and a machine that has truly learned to translate has learned, more or less by accident, to *continue*, and a machine that can continue any text can converse, compose, argue and advise. The Transformer is the latest frame bolted onto the loom this book is named for. Every talking machine now running — every chapter of this history since, every sentence of the flood we live in — is woven on the architecture that eight people built in 2017 to move sentences between German and English a little faster. All eight, I note in passing, have since left Google, mostly to found companies of their own; the firm that owned the loom could not keep the weavers, which any Lyon silk merchant of 1810 could have told them.



Step back now, all the way back to the stone, and look at what the story did — because the ending is folded into the beginning with a neatness that history rarely permits and I refuse to spoil by underselling.

Champollion deciphered Egypt from parallel text and pattern: no phrasebook, no grammar book, no living speaker, the grammar written afterwards as a description of what the counting had already found. That is — precisely, not poetically — how the machines learned. No one gave the Transformer a dictionary. No one told it that French adjectives trail their

nouns or that "pen" has two sizes; it inferred the regularities from alignment and frequency, at a scale of billions of sentences, and our linguistics, like Champollion's grammar, now trails behind the accomplished fact, reverse-engineering what the loom already does. The method of 1822 — align, count, guess, check — turned out to be the method, full stop; Egyptology was the dress rehearsal for the twenty-first century, performed by one man with a pencil. There are Transformers today being aimed at Akkadian cuneiform, at the unread archives of civilizations with no Rosetta Stone of their own; Champollion's heirs have begun handing the pencil to the machine he anticipated. And the priests of Memphis, remember, started all of it with an instruction to copy: carve this decree, in three scripts, in every temple. We obeyed the spirit of that instruction on a planetary scale — we transcribed ourselves into every server on Earth, in every language at once, parliament by parliament, email by email — and out of the copying, exactly as at Rosetta, meaning proved recoverable by pattern alone. The stone was a promise. The loom kept it.

And the ironies, as always, travel in convoy. The money was voted to read the enemy's journals, and the tool ended up writing everyone's letters — the surveillance apparatus of the Cold War matured into the homework aid of every teenager on Earth, which is either swords into ploughshares or something less comfortable, and Chapter Nine will haggle over which. Babel was never repealed: Wilkins's forty boxes and Zamenhof's sixteen rules lie where the twentieth century left them, and the towers still stand — but there is now a basement passage between every pair of them, dug not by scholars or bishops or committees but by counting, and a child in any tower can shout down the stairs and be understood in all the others. That was the entire dream, from Genesis onwards, and we got it — sideways, secondhand, as the exhaust of an arms race and a search engine, which is the only way we ever get anything. And the axemaker's usual invoice arrived stapled to the gift, as it always does: when understanding is free, the

discipline of understanding goes unbought. A generation is growing up that will never wrestle a foreign sentence to the ground — never learn, from the inside, that *pain* is bread and a false friend, that other languages are not costumes over English but other ways of being a mind — because the machine in their pocket has quietly abolished the need. Something real is being paid there, as it was when writing withered memory, and we will not know the size of the bill for fifty years.

But there is one more turn of the screw, and it is the one that leads out of this chapter into the next. The translation story has a moral so strange that the field itself took years to say it aloud: the machine was built to carry meaning *between* languages, and in learning to do that, it had to build — in Weaver's basement, in the geometry of those hundreds of dimensions — something disturbingly like a representation of meaning itself. To translate everything, it had to model everything: boxes, pens, vodka, weather, fisheries, grief. And a model like that cannot be trained on Hansard alone. It needs the biggest parallel text ever assembled — everything mapped against everything, the whole written residue of the species: the encyclopaedias and the love letters, the patents and the poems, the truth and the lies. Which raises a question the priests of Memphis never had to face when they ordered their decree carved in every temple: who owns the temples? The library the machine ate did not assemble itself. It was built over three centuries, by dreamers and philanthropists and volunteers, most of whom gave it away for nothing, none of whom were asked. That story begins with a man in Paris smuggling an encyclopaedia past the censors of the King of France, and it is the next chapter.

CHAPTER EIGHT

Common Knowledge

*The whole human memory can be, and probably in a short time
will be, made accessible to every individual.*

— H. G. WELLS, *WORLD BRAIN* (1938)

Some time in the third century before Christ, if your ship put in at Alexandria — and if you were in the eastern Mediterranean carrying trade, sooner or later it did — you received a visit from the port authorities, and they were not interested in your wine. By standing order of the king, every vessel entering the harbour was searched for books. Any scroll found on board — the story is the physician Galen's, told four centuries later, and he is our only witness, though nobody has ever much wanted to doubt him — was confiscated on the spot and carried up to the great Library, where scribes copied it out fair. And then — here is the detail that tells you everything you need to know about the people involved — the *copy* was handed back to the owner with the king's compliments,

and the *original* went onto the shelves, catalogued, so Galen tells us, under its own special shelf-mark: "from the ships." Think about that for a moment. Not "we will borrow your book." Not even "we will copy your book." We will keep your book, because the original is the good one, and you may have our copy, and you should be grateful, and next time you enter this harbour we will do it again. It was piracy conducted by librarians, with receipts. It was also, though nobody involved could have known it, a rehearsal — the first pass of a pattern that would take two thousand two hundred years to come round again, when machines were sent out into another kind of harbour to read everything that floated past, keep the goods, and offer the world a very fluent copy. The first crawl was conducted with papyrus and gangplanks. Hold the thought; we'll be back at the docks by the end of the chapter.

The Ptolemies, to be fair, were collectors on a scale that makes the behaviour almost forgivable. Ptolemy III, Galen tells us in the same breath, wrote to the city of Athens asking, very politely, to borrow the official state copies of Aeschylus, Sophocles and Euripides — the master texts, kept in the public record office, against which all others were corrected — so that his scribes might make a transcription. Athens, having been burned before (possibly literally), demanded a deposit of fifteen talents, a sum you could have run a small war on. Ptolemy paid it, took the scrolls home, kept the originals, sent Athens the copies, and told them to keep the money. He regarded it, we may assume, as a bargain, and by the standards of what those texts have since become, he was right. The ambition behind all this was stated quite openly: the Library was to hold everything — every book of every people, translated into Greek, gathered in one place, so that (the phrase belongs to the whole enterprise if not to any single Ptolemy) a man standing in one room might consult the entire memory of the species. Ancient authors claim half a million rolls, and ancient authors, it must be said, claim all sorts of things; the honest answer is that nobody knows, and that it hardly matters, because the number was never the point. The point

was the appetite. The Library of Alexandria was the first institution in history founded on the proposition that knowledge should be *complete* — and, note carefully, complete in one particular building, belonging to one particular king. Universal access for everybody who could get to Alexandria and past the door staff. The commons, with a landlord.

And then it burned, as everyone knows. Except that it didn't — or rather, it did, repeatedly, at the hands of almost everybody, which is the surest sign in history that a story is doing work as a myth rather than as a fact. Julius Caesar is blamed, because in 48 BC his troops fired the ships in the harbour and the flames spread to dockside warehouses that seem to have been full of books awaiting cataloguing — from the ships, one assumes, to the last — yet the geographer Strabo was working contentedly in the Library's district a generation later. The Christians are blamed, because in AD 391 a mob under the bishop Theophilus demolished the temple of Serapis, which housed the daughter library; that one has at least a decent claim to being partly true. And the Muslims are blamed, thanks to the story that the Caliph Omar, asked in 641 what should be done with the books of Alexandria, replied that if they agreed with the Koran they were superfluous and if they disagreed they were pernicious, and either way should heat the city's bathhouses, which they allegedly did, for six months. It is a marvellous line. It is also first recorded nearly six centuries after the event, by writers with excellent reasons for telling tales about somebody else's caliph, and almost no serious historian believes a word of it. I pass it on as I pass on the story about the chairman of IBM and the five computers (Chapter Eleven has it in stock): flagged, winking, because we need it. What actually killed the Library, so far as anyone can tell, was nothing so cinematic. The money dried up. The kings and then the emperors lost interest; one of the later Ptolemies purged the scholars in a political tantrum and the best of them scattered to Rhodes and Pergamum; and papyrus, which rots, survives only so long as somebody is paid to keep recopying it. A library is not a building full of books. It is a *subscription* —

a standing agreement, renewed every generation, to keep paying the copyists — and when the subscription lapses, the collection dies of neglect long before anyone bothers to burn it. Remember that, too. Nobody ever destroys the great library. They just stop funding it, and then argue for two thousand years about whose fire it was.

This chapter is about that library — not the building, the *project*: the long, lurching, repeatedly bankrupted attempt to gather up everything everybody knows and give it to everybody else. It is the chapter in which the thread of this book's title actually gets spun. Everything we ever said: the corpus the talking machine was woven from did not fall from the sky; it was accumulated, scroll by scroll and card by card and edit by edit, by twenty-three centuries of collectors, encyclopedists, robber barons, dreamers, librarians and unpaid volunteers, most of whom believed — it is the one conviction they all share, from Ptolemy's scribes to a teenager correcting a comma in Wikipedia at two in the morning — that knowledge gathered belongs, somehow, to everyone. And it is the chapter in which we ask the question that conviction has always dodged and can now dodge no longer: *does it?* Because the whole gathered hoard has just been read, end to end, by a machine that turned it into the most valuable property on Earth, and the courts are full, and the question of who owns what everybody wrote has stopped being philosophy and started being money.



Owning knowledge used to be beautifully simple, because knowledge was an object, and objects can be chained to the furniture. Visit Hereford Cathedral and you can still see the arrangement: the largest surviving chained library in the world, seventeenth-century oak presses, and hanging from the fore-edge of each volume a metre or so of iron chain running on a rod, long enough to let you lift the book to the reading desk, short enough to stop you leaving with it. This was not paranoia; it was

arithmetic. A book, in the age before print — the age whose collapse Chapter Eleven anatomizes, so one sentence here will do: this is the pattern this book keeps returning to — cost as much as a farm, and you do not leave a farm lying on a desk. The Sorbonne chained its reference collection in the 1290s; Oxford and Cambridge chained theirs for four hundred years; and the chains came off, library by library, in almost perfect synchrony with the falling price of the printed page, until by the eighteenth century the ironmongery survived only in backwaters like Hereford, as a curiosity. I ask you to look hard at that chain, because it is the most honest piece of information technology ever built. It says: *you may read this, and you may not take it away*. Every licence agreement you have ever scrolled past without reading, every paywall, every lawyer's letter about training data, is that chain, reformed in language. The library has always had two souls — come in, says one; hands off, says the other — and the entire history in this chapter is the two of them arguing.

By the seventeenth century the argument had produced something genuinely new: a community that treated knowledge as common property across borders that armies couldn't cross. They called it the Republic of Letters — a republic with no territory, no government and no taxes, whose citizens were anybody with Latin, a pen and the postage: Boyle writing to Huygens, Leibniz writing to everybody, thousands of letters a year moving between men (and a scandalously small number of women) who often never met, on the explicit understanding that a discovery hoarded was a discovery wasted. It was snobbish, quarrelsome and gloriously leaky, and it established the precedent everything since has traded on: that finding something out creates a debt to everyone else who is trying to.

And in 1731 a Philadelphia printer worked out how to let ordinary tradesmen buy citizenship in that republic by subscription. Benjamin Franklin and his Junto — a self-improvement club of clerks, surveyors and shopkeepers who could not individually afford books — pooled their money: fifty subscribers, forty shillings each, ten shillings a year thereafter,

the fund to buy books held in common, each subscriber owning a share. The Library Company of Philadelphia, Franklin wrote later with pardonable smugness, was the mother of all the North American subscription libraries, and he credited it with making American tradesmen as well-read as gentlemen elsewhere. But notice the mechanism, because we shall want it again, much later, when the stakes have grown beyond anything Franklin would credit: the common stock of knowledge was owned by *shareholders* — everyone who had paid in held a piece of the whole, and nobody owned any book in particular. It is, so far as I know, the first working answer to this chapter's question. Who owns the library everybody built? Everybody. In writing. With certificates.



Now to Paris, and the most dangerous reference work ever published.

In February 1752 the King's Council of State suppressed the first two volumes of a new encyclopedia, on the grounds that it was destroying royal authority and corrupting morals, which, to be fair, it was. Word reached its editor, Denis Diderot, that a police raid was coming for his manuscripts — years of work, thousands of articles in draft. He needed a hiding place no officer of the Crown would ever search. He found one. The manuscripts went, at the owner's own suggestion, to the house of Chrétien-Guillaume de Lamoignon de Malesherbes — the Director of the Book Trade, the royal official personally in charge of censoring the *Encyclopédie* — who reasoned, correctly, that nobody raids the censor. The story comes to us from Diderot's daughter, and if it has been polished in the telling, it deserved the polish. It is also, once you stop laughing, the key to the whole affair: the *Encyclopédie* survived because a significant fraction of the French establishment, up to and including the woman who shared the king's bed, wanted it to survive, and merely required that the forms of suppression be observed. Enlightenment France banned the book with one hand and subscribed with the other, which tells you the *ancien régime* was

already hollow; it took another thirty-seven years for anyone to knock on it and hear the echo.

The thing had begun, as most revolutions do, as a modest commercial proposition. A Paris publisher, André Le Breton, wanted a French translation of Ephraim Chambers's two-volume English *Cyclopaedia*, and hired as editors a young mathematician of European reputation, Jean d'Alembert — who had been abandoned as a baby on the steps of a Paris church, a foundling who grew up to co-edit the sum of human knowledge, and no novelist would dare — and a jobbing man of letters named Diderot, conveniently both brilliant and cheap, who would shortly interrupt the work with a spell in the fortress of Vincennes for an impious essay. In their hands the translation swelled into something else entirely: seventeen volumes of text, eleven of plates, some seventy thousand articles by a hundred and forty contributors, a quarter of a century in the making. Diderot announced the ambition in the article "Encyclopédie" itself, and it is worth quoting in spirit if not in letter: the purpose of an encyclopedia, he declared, was to gather the knowledge scattered across the face of the earth, to set out its general system to the men with whom we live, and to transmit it to those who come after — so that our descendants, being better instructed, might be happier and more virtuous. And, he added, in the phrase that got him into trouble and belongs on the title page of this book: *to change the common way of thinking*.

Look at what that meant in practice, because it connects to everything Chapter Eleven says about mystery and margin. The Encyclopédie's most subversive contents were not the articles on God, though those caused the noise. They were the plates: eleven volumes of exquisite engravings showing, in exploded diagrams and numbered figures, exactly how things were made — pin-making, glass-blowing, type-founding, the stocking-frame, tanning, gunpowder. For centuries the guilds had guarded their crafts as "mysteries," a legal term, with apprenticeships and oaths to match; Diderot, a cutler's son, sent artists into the workshops to draw it all, and

sold the mysteries by subscription to anyone with the price. It was the expertise-collapse of Chapter Eleven, run at one-twentieth speed on copper plates. And where the censors did force discretion, the editors invented a weapon so quietly perfect that I want you to notice its shape: the cross-reference. An article on some doctrine of the Church would be a model of orthodox piety — and carry at its foot a small italic instruction: *see SUPERSTITION*. The infamous example is the entry on cannibals, "Anthrophages," which concludes, deadpan, by directing the curious reader to the articles on the Eucharist and Communion. The censor read articles; the sedition lived in the links. File that away: two hundred years before anyone coined the word, the most explosive publication of the eighteenth century was already hypertext, and already demonstrating the principle on which the web, Wikipedia and the talking machine would all eventually run — that the meaning is not in the entries. It is in the connections. You may imagine how much I enjoy that.

The rest of the story is a war of attrition, and I will give you only the good parts. In 1757 a madman scratched Louis XV with a penknife, the government panicked about impious literature, and the screws tightened; in 1759 the Council revoked the licence outright and the Pope placed the work on the Index. D'Alembert, who had had enough (an article on Geneva had enraged Calvinists and Catholics simultaneously, a genuine achievement), resigned the co-editorship, leaving Diderot to carry the enterprise more or less alone — alone, that is, except for the Chevalier de Jaucourt, a Protestant nobleman-physician who wrote, unpaid, roughly a quarter of the entire *Encyclopédie*, article after article, day after day, for fifteen years, employing a small office of secretaries at his own expense; the story goes that he sold a house to pay them, and sold it, with perfect circularity, to Le Breton the publisher. Remember Jaucourt whenever anyone tells you nobody works for nothing: the largest single contribution to the eighteenth century's greatest commercial publication was donated, by a man history nearly forgot. He is the patron saint of a certain website

we shall reach presently. The banned volumes, meanwhile, were printed anyway — all ten remaining text volumes at once, in 1765, under a false imprint claiming they came from Neuchâtel in Switzerland, a fiction precisely nobody believed and the authorities, Malesherbes's fingerprints all over the arrangement, declined to examine. And then, the exquisite final twist: in 1764 Diderot discovered that the real censor had been inside the house all along. Le Breton, terrified for his investment, had for years been secretly cutting the most dangerous passages from the proofs *after* Diderot signed them off, quietly, at the print shop, where no one would ever know. Diderot wrote him a letter of magnificent fury — he had massacred, Diderot told him, the work of twenty good men, and the wound would never heal — and then, because there was nothing else to do, went back to work for him. So tally the affair: the state's censor protected the book, and the book's own publisher censored it. The men paid to suppress knowledge smuggled it, and the man paid to sell it mutilated it. If you want to understand everything that follows in this chapter — Minnesota's licence and CERN's gift, the shadow libraries, the lawsuits — you may take the *Encyclopédie* as your Rosetta Stone: guardianship and ownership are different jobs, they are frequently held by the wrong people, and the commons has never once been able to tell, at the time, which of its friends were which. By 1789, counting the cheap pirated editions that poured across every border in Europe, perhaps twenty-five thousand sets were in circulation. What the readers did with them, four years into the Revolution, rather suggests the common way of thinking had changed on schedule.



The nineteenth century's great contribution to the library was to work out who should pay for it, and the answer it produced was: a man with a guilty conscience and a great deal of steel.

In 1848, in Allegheny, Pennsylvania, a Scottish immigrant boy was working in a cotton mill as a bobbin boy for a dollar twenty a week, twelve hours a day, and a local worthy named Colonel James Anderson opened his private library — four hundred volumes — to working boys, free, on Saturday afternoons. The boy borrowed a book a week, and never got over it; sixty years later he could still recite the librarian's hours. When the library's keeper decided the free privilege applied only to boys formally bound as apprentices, and not to mere wage-workers like him, the boy wrote an indignant letter to the *Pittsburgh Dispatch* signed "Working Boy," and won, which may have been the most consequential letter to the editor ever published, because the boy was Andrew Carnegie, and when he became the richest man in the world he remembered exactly what a shelf of borrowed books feels like from below.

Carnegie's essay "Wealth" — the world quickly renamed it "The Gospel of Wealth," and he didn't object — appeared in 1889 and said something no plutocrat had ever said in print: that a rich man's fortune was not really his; that he was merely its trustee for the community that had made it possible; and that "the man who dies thus rich dies disgraced." He then spent the rest of his life executing the doctrine with the same ferocious system he had applied to steel. Two thousand five hundred and nine libraries — nearly seventeen hundred of them across the United States, the rest from Dunfermline to Fiji — built to a formula as unsentimental as a rail contract: Carnegie supplied the building; the town had to supply the site, and bind itself in writing to tax itself ten per cent of the grant, every year, forever, for books, staff and heat. He had grasped the lesson of Alexandria without, one assumes, ever putting it that way: a library is not a building, it is a subscription, and a philanthropist who funds the building without securing the subscription has bought the town a handsome ruin on a thirty-year delay. The formula worked. For a couple of generations the Carnegie library — free to all, and the words were frequently carved over

the door — was the closest thing the English-speaking working class had to a university.

And now the wryness, which the reader has been expecting. The money for the temples of knowledge came out of the mills, and in 1892, at Homestead, just up the Monongahela from a town Carnegie had already given a library, his company broke a strike with three hundred armed Pinkerton agents brought up the river on barges; the battle that followed killed men on both sides, the union was destroyed, and the twelve-hour shift rolled on. Carnegie was in Scotland at the time, trout-fishing, having left instructions of studied ambiguity, and his critics observed — they observed it loudly, at the time, in print — that a steelworker coming off a twelve-hour turn had precious little daylight left for the reading room his wages had subsidized, and that what workers wanted was not free books but a living wage and eight hours to enjoy it in; several dozen communities refused Carnegie's money outright as tainted. They had a point, and I do not propose to take it away from them. Steel was laundered into knowledge; the Gospel of Wealth was, among other things, the world's most successful reputation-cleaning operation. And also: the boy at Colonel Anderson's shelf was real, and so were the millions of other children who walked up those library steps for a century afterwards, and both books of the ledger are true at once, which is what makes history history instead of a morality play. The robber barons built the commons with one hand and picked its pocket with the other. Hold that thought too. It is about to become extremely contemporary.



Meanwhile in Brussels, a lawyer had decided that books themselves were the problem.

Paul Otlet — say it *Ot-LAY*, he was Belgian and has suffered enough — noticed something in the 1890s that every user of the world's libraries had felt and nobody had said: the book is a terrible container. The fact you

need lies buried in a chapter, in a volume, on a shelf, in a city you are not in. What if you sliced knowledge finer — down to the individual *fact* — and put each one on its own index card, classified by a universal decimal code so that every card about, say, aerial navigation sat with every other, regardless of what book, language or country it had arrived from? Otlet and his partner Henri La Fontaine (who in his spare time won the Nobel Peace Prize, as one did in that optimistic hour) founded their International Institute of Bibliography in 1895 and started carding the world's literature. By the 1930s the collection — housed by then in a wing of a government palace in Brussels, under the deliriously wonderful name of the Mundaneum — had passed twelve million cards, and had become, quite literally, a search engine: anyone on Earth could post or wire a query to Brussels, and for a fee the staff would run the question through the drawers and mail back copies of the relevant cards. Some fifteen hundred queries a year, answered by hand, at walking pace. And Otlet, who could see perfectly well that paper was only the larval stage, wrote down where it was heading: work desks fitted with screens; documents summoned from a distance by wire; sound and moving images joined to text; the whole apparatus converging, in his words, into a mechanical, collective brain. He published that in the mid-1930s. In 1934 the Belgian government, bored with the whole affair, shut the palace wing; in 1940 the Germans arrived in Brussels and cleared out tons of the collection — cards, pamphlets, posters, decades of it — to make room for an exhibition of Third Reich art. Otlet died in 1944, ruined and obscure, and stayed obscure for forty years until historians dug him up and the internet's own children recognized, with a shiver, their grandfather. There is a Mundaneum museum in Mons now, keeping what survived of the drawers. Its sponsors, for some years, included a certain large search-engine company, which is either touching or grave-robbery depending on the light, and I honestly cannot decide, and that ambivalence, reader, is the entire subject of this chapter in miniature.

H. G. Wells took the same dream and did what Wells always did with a dream: gave lectures about it. Through 1936, 1937 and 1938 — at the Royal Institution, in America, anywhere that would have him — he preached what he called the World Brain: a permanent world encyclopedia, continuously revised by an organized corps of the world's knowledge-workers, reproduced by the miraculous new microfilm until any student anywhere could consult an exact replica of the whole, and standing above the squabbling nations as a common ground of verified fact. The sentence at the head of this chapter is his, and note the confidence of it: the whole human memory, accessible to every individual, *probably in a short time*. He believed — this is the part that dates him, and the part I find hardest to mock — that such an encyclopedia would not merely inform the world but *pacify* it: that war grew in the gaps between what men knew, and a common brain would close the gaps. Within a year of the book's publication the most literate nations on Earth were at war with each other, and Wells lived just long enough to see the newsreels from the camps and die in 1946, out of patience with the species. On the pacifying power of shared information, the twentieth century returned a verdict, and the twenty-first has co-signed it. On the accessibility, though — on the plain engineering claim — old Wells was dead right, give or take ninety years, and we are inside his sentence now.

It was left to an American engineer to draw the blueprints. In July 1945, in the *Atlantic Monthly*, Vannevar Bush — the man who had run America's wartime science, six thousand researchers, radar to the Bomb — published an essay called "As We May Think," asking what all that organized ingenuity should do next, and answering: fix the library. His imagined machine, the memex, was a desk: microfilm inside, screens on top, and the crucial idea — the reader could tie any two documents together and keep the ties, building trails of association through the record the way the mind itself hops from thing to thing, and could pass whole trails to a friend. "Wholly new forms of encyclopedias will appear," he

wrote, "ready made with a mesh of associative trails running through them." The memex was never built; it didn't need to be. The essay did the building. A twenty-year-old radar technician named Douglas Engelbart read it in a Red Cross library in the Philippines, waiting for his ship home, and spent the next forty years constructing what it described — the screen, the pointer, the link — and through Engelbart the trail runs on, as trails do, to everything you clicked today. Diderot's *renvois*, Otlet's cards, Bush's trails: three times in two centuries the same discovery, that knowledge is not a pile but a web, and that whoever holds the links holds the thing itself.



Which brings us to the two men who actually built it, and the fork in the road between them, and the fork is the hinge of this whole chapter, so slow down and take it carefully.

The first is Ted Nelson, who in 1960, as a Harvard graduate student, began a term project that has now been running for sixty-five years. Nelson is the man who coined the word "hypertext," in 1965, and who saw everything early: that documents would live on screens, that they would quote and embed and link each other endlessly ("everything," he wrote, in the technical vocabulary he built himself, "is deeply intertwined"). His system, Xanadu, was designed with a rigour the web has never had: links that ran both ways, so a document always knew who was quoting it; no copies, but "transclusion" — the quoted passage living once, in the original, and appearing everywhere it was cited; and, at the heart of the design, *royalties*: every time your words appeared on someone's screen, a micropayment flowed back to you, automatically, because the system never lost track of whose words they were. Read that again. Thirty years before the web and sixty before the lawsuits, Ted Nelson designed the answer to this chapter's question — who owns what everybody wrote? *everybody, traceably, at a fraction of a cent per reading* — and built it into the archi-

ecture. And Xanadu never shipped. It became the industry's tragicomic legend; a famous magazine profile in 1995 called it the longest-running vapourware in the history of computing, and Nelson has never forgiven them, and a working demonstration did finally appear in 2014, fifty-four years in, to affectionate applause. I refuse to laugh at him. The reason Xanadu never shipped is precisely that it tried to keep the ledger — every link two-way, every quotation accounted, every author paid — and keeping the ledger of who-owns-what across all human utterance turns out to be, roughly speaking, the hardest problem in the world. Nelson insisted on solving ownership before shipping. Remember that as you watch what shipped.

The second man is Tim Berners-Lee, a software engineer at CERN, the European particle-physics laboratory outside Geneva, who in March 1989 submitted to his boss a memo with the numbing title "Information Management: A Proposal," suggesting a web of linked documents to stop the laboratory losing track of itself. His boss, Mike Sendall, pencilled on the cover the most consequential margin note since Fermat: "Vague, but exciting." Given a NeXT computer to play with, Berners-Lee had the thing running by Christmas 1990 — browser, server, the first website, and a hand-written sticker on the machine reading "This machine is a server. DO NOT POWER DOWN!!" — and in the summer of 1991 he announced it, modestly, in an internet discussion group, as one more hypertext system among several. Technically, it was the crude one. Where Xanadu kept the ledger, the web shrugged: links ran one way only; a document had no idea who linked to it; nothing was accounted, nothing was paid, and if the thing you pointed at vanished, you got an error — the famous 404 — and the world moved on. Every academic hypertext expert could tell you these were defects. They were the entire success. Because a system with no ledger has no gatekeeper, and no gatekeeper means nobody's permission to ask: anyone could link to anyone, claim a page, join the weave, and the broken links could lie where they fell. The web is the Republic of Letters with the

postage abolished — scruffy, unaccounted, gloriously leaky — and it grew like nothing human beings have ever built, precisely because nobody was keeping track of who owed what to whom.

And then, in the spring of 1993, the fork itself — as clean a controlled experiment as history ever runs. The web had a rival, and the rival was winning: Gopher, out of the University of Minnesota, a tidy menu-driven system for serving documents across the internet, simpler than the web and at that moment more widely used. In the spring of 1993 the university announced that commercial users of the Gopher server software would owe a licensing fee. That was all. A reasonable, modest, defensible assertion of ownership — somebody built this, somebody should pay. And the effect, Berners-Lee recorded, was immediate and total: industry dropped Gopher like a hot brick, no company daring to build on ground that had just shown it could charge rent. Watching this, Berners-Lee and his colleague Robert Cailliau pressed CERN to do the opposite, formally and forever, and on the 30th of April 1993 CERN signed a declaration placing the web's protocols and code in the public domain, irrevocably, royalty-free, for anyone, for any purpose. A deed of gift, one page. It may be the single most valuable document of the twentieth century, and the value flowed, note well, from its renunciation: the web beat Gopher because the web was the one you could not be charged rent on. Sixty years of dreaming — Otlet's brain, Wells's encyclopedia, Bush's trails, Nelson's ledger — and when the world's library finally got built, it got built on the version that had *dropped the question of ownership on the floor*. The commons won because it stopped keeping accounts. The bill, of course, does not vanish just because no one is keeping it. It goes into a drawer, and compounds, and thirty years later something with no face pulls the drawer open. But I am getting ahead of myself.

★

First, the encyclopedia — because the web, having conquered, was a library with no reference desk, and the old dream demanded its due.

In March 2000 an internet entrepreneur named Jimmy Wales and a philosophy doctoral student named Larry Sanger launched Nupedia: a free online encyclopedia, written by credentialed experts, every article to pass a seven-stage gauntlet of review and revision, rigour throughout. It was Wells's World Brain as Wells had specified it — an organized corps of the qualified — and it worked exactly as well as a committee of the qualified always works: after roughly a year, about a dozen finished articles. The whole human memory, at that rate, was scheduled for completion some time after the heat death of the universe. So in January 2001, as a side project, almost an apology, they bolted on a scratchpad where anybody could draft material for the experts to review later, using a piece of software Ward Cunningham had invented in 1995 and named after the shuttle bus at Honolulu airport — *wiki wiki*, Hawaiian for "quick." You know what happened, but dwell on how completely nobody else did. The scratchpad produced more in two weeks than the encyclopedia had in a year. It ate its parent, and then it ate the world: a thousand articles within months, twenty thousand by year's end, millions within the decade, in hundreds of languages, written by nobody in particular, owned by nobody at all, free to everyone, forever. Every sane prediction said it must fail — and the predictions were not stupid; they were merely arguing from all previous human experience. Anyone can edit meant any fool or liar can edit, and fools and liars are legion, and yet the thing converged: it turned out — mostly, on average, with famous and instructive exceptions — that the people who care to fix an article outnumber and outlast the people who care to wreck it, that a page watched by ten strangers is guarded better than a vault, and when the journal *Nature* sent expert reviewers through matched science articles in 2005, they found the upstart and the *Encyclopaedia Britannica* in roughly the same league for errors — the amateurs slightly scruffier, nothing more. Wikipedia works in practice, its editors

like to say, not in theory. Seven years after that study, Britannica ended its print edition, two hundred and forty-four years after Edinburgh, and the mystery went the way mysteries go in this book. And presiding over the whole improbability, if you look with the right kind of eyes, is a familiar ghost: the Chevalier de Jaucourt at his desk, unpaid, uncredited, eight articles a day for the common stock, multiplied now by a hundred thousand. Diderot needed a hundred and forty contributors and got one saint. Wikipedia found the saints were a standing crop of the species, waiting only for a door with no lock on it. Strangers can be trusted, mostly. It is, when you set it against everything else your newspaper carries, among the most hopeful facts ever demonstrated.



The gathering hoard kept two sets of shelves, though, and the other set stayed chained. The world's scientific literature — the Republic of Letters' own descendants, papers largely paid for from public purses, written free, reviewed free — sat behind subscription tolls priced for institutions, so that the farmer's son in Nairobi or, for that matter, the curious taxpayer in Ohio who had funded the work could see the abstract and pay forty dollars for the rest. One of the people who found this arrangement intolerable was a young American named Aaron Swartz, and he had standing to mind: at fourteen he had helped write the standard by which the web syndicated its news; he had helped build the machinery of Creative Commons, the licences by which authors give their work to the commons deliberately; he was, by every account, one of those people who treat the axemaker's gift as a public trust. "Information is power," he wrote in 2008. "But like all power, there are those who want to keep it for themselves." In 2010 and 2011, using a laptop wired into a network closet at MIT, he downloaded several million articles from the scholarly archive JSTOR — to what ultimate purpose was never established, because he never got the chance to say. JSTOR got its files back and declined to pursue him. The federal prosecutors did not

decline: the charges, as stacked, carried a theoretical maximum of decades in prison, for downloading academic papers. In January 2013, with the case still pending, Aaron Swartz took his own life. He was twenty-six. I will not put a moral in his mouth; he paid too much for the privilege of his own. I will only set the fact down next to the others in this chapter and let the shelf do the arguing: Diderot went to Vincennes, Tyndale will go to the stake in Chapter Eleven, and the twenty-first century, which mints a fortune every hour from reading everything everyone ever wrote, hounded a young man to death over a closet full of journal articles. The chain library never went away. It just got better lawyers.



And so, the docks again.

While the courts were still occupied with Swartz, the machines had already begun eating. It started respectably: in 2004 a search-engine company undertook to scan the world's books — tens of millions of volumes, whole university libraries fed page by page through the cameras — to make them searchable, and the authors' representatives sued, and after a decade the American courts ruled the scanning lawful: fair use, said the judges, because an index that tells you a book *contains* your answer is a new and transformative thing, not a rival copy; the machine had read the books in order to point at them. That precedent — reading in order to point is fair — held the door. And through the held door, in the years that followed, came something with a rather larger appetite. A small nonprofit called Common Crawl had been quietly doing, since 2007, exactly what Ptolemy's harbourmasters did: sweeping the open web, page by page, billions of pages, and storing the haul — but in the Franklin spirit, not the Ptolemy one, giving the copies away free so that researchers without a giant's resources could study the web whole. Petabytes of everything we ever said, boxed and public: the commons, archived by idealists, as a gift. And when the builders of the talking machines needed thread for the loom,

there it lay. The crawl went in. The books went in — the scanned, the licensed, and, as later filings would air in public, collections of rather murkier provenance that had circulated for years in the internet's shadow stacks. The encyclopedia that anyone can edit went in, all of it, the single richest seam in the weave. The forums went in, the reviews, the recipes, the code, the love letters posted where crawlers could reach them, your paragraphs and probably mine. Alexandria took two and a half centuries to gather half a million scrolls and never read most of them. The machine read the equivalent of the whole human library — read it in the only sense that ends up mattering, took the pattern of it into itself — in a matter of weeks, on hardware described in Chapter Five, by the method described in Chapter Six. And this time, notice, the harbourmaster's etiquette had decayed. Ptolemy kept your original but at least handed you a copy, and his catalogue said honestly *from the ships*. The loom keeps no shelf-marks. The thread goes in labelled and comes out woven, and you cannot point to the cloth and say: that inch is mine.

Then the librarians sued — and the authors, and the newspapers, and the photographers, each in their own actions, and I shall describe the class of dispute rather than the docket, partly because the cases were still crawling through the courts as I wrote and partly because the docket is not where the interest lies. The plaintiffs' claim is the oldest in this chapter and you have heard it at Hereford: *you may read, and you may not take away* — our archives were made by paid reporters and hungry novelists, they were never anybody's to weave, and the weave now answers, fluently, the very questions our work exists to sell. The defence is nearly as old and you have heard it too: the machine keeps no copies; it did to the library only what every reader does, which is to be changed by it; we do not sue the young writer for having read her predecessors, and everything in the commons was built by exactly that kind of borrowing, all the way down. The maddening thing — I have been on both sides of this argument in a single afternoon, and so, if you are honest, have you — is that each side is

holding a genuine piece. Scale is real: a reader absorbs a library in a lifetime and a machine absorbs it before lunch, and the law has never before had to decide whether a difference of ten million times is a difference of kind. And transformation is real too: the cloth genuinely is not the thread, and a doctrine built for quotation and parody is being asked to govern digestion. The courts will split differences, as courts do; settlements will draw the first crude price-tags on the previously priceless; the outcomes will vary by country and decade and will keep this generation of lawyers in the manner to which the next chapter's guilds were accustomed. All that is weather. Here is the climate.

The commons was built by everyone, for everyone. That was not incidental to how it was built; it was the engine. Diderot's subscribers, Franklin's shareholders, Carnegie's readers taxing themselves ten per cent a year, Jaucourt and his hundred thousand descendants donating the articles, Berners-Lee signing away the rent, the idealists boxing up the crawl for science, you, every time you answered a stranger's question in a forum at midnight because someone once answered yours: the whole hoard was assembled on the explicit or implicit promise that it was *for everybody*, which is precisely why no one guarded it. And the loom wove that unguarded everything into the most valuable artefact of the age — valuable *because* the thread was everything, and universal *because* the commons was — and the artefact is private property, and access to it is metered, and the rent is collected by the handful of institutions rich enough to own the looms, whom you will meet at the end of Chapter Eleven, still holding the handle of the axe. The robber baron laundered steel into libraries; this runs the wash the other way. So the question — who owns the pattern woven from everybody's thread? — will not be settled by the lawsuits, because the lawsuits can only ask it author by author, archive by archive, and the true answer is not divisible that way. Your sentence in the weave is unfindable, untraceable, worth nothing on its own and irreplaceable in aggregate, exactly like your vote. Nelson's ledger, which would have tracked every

thread to its spinner, never shipped, and can never ship now; the accounting was skipped at the moment of weaving and cannot be reconstructed. Which leaves, as far as I can see, only one honest direction to look, and it is back towards a Philadelphia printer and his forty-shilling shares. If no one's thread can be found in the cloth, but the cloth is made of everyone's thread, then the claim that survives is not payment for your contribution. It is a *share in the whole* — the Library Company's answer, scaled to the species: everyone who paid into the common stock owns a piece of the common stock, and the machine woven from everything everyone ever made owes something like a dividend to everyone. What that could possibly mean in practice — what it does to money and work and the state when the payout is not a metaphor — is a question I am not yet equipped to open, because it belongs to the far end of this book, where scarcity itself comes up for renegotiation. Chapter Twelve will try. Consider the seed planted, and the ground, I hope, disturbed.

For now, stand a last moment on the dock at Alexandria, because both futures were already berthed there. A king who kept the originals, and a library that meant to hold everything for everyone; the grasp and the gift, in the same building, run by the same people, and twenty-three centuries in which neither ever managed to finish the other off. Wells promised the whole human memory to every individual, probably in a short time, and it took ninety years and arrived — freely, at three in the morning, in every language, exactly as dreamed — owned. The chain and the open door, both at once, which is where this chapter came in. History says the chains come off; it has said so at Hereford and Mainz and Minnesota and Geneva, every single time, eventually. History also says — Chapter Eleven will teach you the ugly word — that *eventually* sends no schedule. The scrolls, at least, are finally off the ships. What we are owed for them is the next argument, and, I rather think, the great one; and while the lawyers warm up, the machine that read everybody's work has quietly started doing

everybody's job — which is a different chapter, and it begins on the next page.

CHAPTER NINE

The New Apprentice

We can know more than we can tell.

— MICHAEL POLANYI, *THE TACIT DIMENSION*

In 1888, in Kansas City, Missouri, an undertaker named Almon Brown Strowger was losing business, and he believed he knew exactly where it was going. In those days you did not dial a telephone; you lifted the earpiece and asked a young woman at a switchboard to connect you, and the young woman on the Kansas City switchboard — so the story goes, and I pass it on at the speed of gossip, because gossip is the only vehicle it has ever travelled in — was married to a rival undertaker. When the bereaved of Kansas City telephoned to ask for Strowger, Strowger became convinced, the call was quietly diverted to the competition, and the competition got the corpse. Now, a man who suspects the telephone operator of stealing his funerals has several options. He can complain to the telephone company. He can hire a lawyer. Strowger, bless him, chose the

third option: he set about abolishing the telephone operator, as a category of human being, from the face of the earth. Using a round collar box and some straight pins, he worked out a switch that would let the caller select a number for himself, by electrical pulses, with no young woman anywhere in the circuit; he filed his patent in 1889 and had it granted two years later; the first automatic exchange opened in La Porte, Indiana, in 1892; and his company advertised the result, with the delicacy of the age, as the "girl-less, cuss-less, out-of-order-less, wait-less telephone." Descendants of Strowger's switch were still clattering away in exchanges around the world nearly a century later. It is possibly the pettiest origin story in the history of engineering — one of the great labour-displacing machines of the twentieth century, conceived not to lower costs or raise throughput but to stop one particular woman doing one particular thing — and I begin with it because it contains, in miniature, everything this chapter is about. A skilled human being stood in the middle of a system. The skill was real: a good operator held the geography of a city's conversations in her head, recognized voices, routed emergencies, covered for the machine's failures. And the skill made her powerful, and the power made someone distrust her, and the distrust built a machine, and the machine ate the job. Multiply Strowger's collar box by everything, and you have the industrial history of the last two centuries.

But this chapter is not, mainly, about jobs. The economics of what the talking machine does to the pay-packet — who is displaced, what softens the landing, what the whole upheaval is for — will get its hearing in the closing chapters, and history's answers are less bleak and less comfortable than you expect. This chapter is about something underneath the jobs, something older and, I think, more fragile: the craft itself. Skill — how a human being comes to have it, how it passes from one pair of hands to another, and what happens to the chain of transmission, which is far and away the oldest institution our species possesses, when the newest link in the chain is a machine. Because that is what has actually happened, and

almost nobody has put it this way: we have not built a tool, in the sense that a hammer is a tool. We have taken on an apprentice. It arrived knowing nothing, it learned by watching us — by watching, quite literally, everything we ever wrote, drew, coded and said — and it has now finished its indentures and is asking, politely, at any hour you please, about the terms of its mastership. The guilds had rules for this moment. We have forgotten them. It may be worth remembering what they were.



Begin with the word *indenture*, because the word is a picture. When a fourteen-year-old boy was bound apprentice in medieval England, the contract between his father and the master was written out twice on a single sheet of parchment, which was then cut in two along a deliberately jagged, toothed line — *indented*, from the Latin for teeth — so that in any future dispute the two halves could be fitted together and proven genuine. No seal, no notary, no central registry: the document authenticated itself, by the impossibility of forging a torn edge. Keep that image by you; we shall want it again at the end of the chapter, when tearing an unforgeable edge becomes, once more, the problem of the age.

The terms inside the indenture were ferocious. The boy would live in his master's house for seven years, be fed and clothed and taught the craft, and in exchange would serve utterly: the standard London forms bind him not to marry, not to fornicate, not to play at dice or cards, not to haunt taverns or playhouses, not to absent himself night or day, and not — a clause with a certain poignancy — to reveal the secrets of his master's trade. Seven years, for a craft that a bright boy could largely pick up in three; the cynical reading of those extra years — that they protected the scarcity of masters rather than the standards of the craft — is true, and will get its full airing later in this book, so let the guilds off with a caution here, because the cynical reading is not the whole truth. Something else was being transmitted in those seven years, something that genuinely could not be

hurried, and to see what it was you have to follow the boy out the far end of his term. He did not become a master. He became a *journeyman* — from the French *journée*, a man paid by the day — and in much of Europe he was now expected to leave. Not metaphorically: to walk out of the town gate and keep walking. The German-speaking lands turned this into a formal institution, the *Wanderschaft* — the wandering years — under rules of almost liturgical precision: the young journeyman travelled for three years and a day; he was forbidden to come within a wide ring of his home town for the duration; he carried a book in which the masters he worked for entered his conduct; and he presented himself in each new town at the appointed inn, worked weeks or months in a strange workshop, learned how *they* did it — the different clay, the different grain, the different trick with the plane — and moved on. You may, if you keep your eyes open in southern Germany today, still occasionally see one: a young carpenter in the traditional black corduroys and broad hat, thumbing a lift with a walking-stick and a bundle, because the institution, astonishingly, is not quite dead. A few hundred still go on the *Walz* every year, eight centuries after the custom began.

And then, at the end of the wandering, if he wished to be a master, the journeyman submitted to the guild a single piece of work into which he was to put everything he had — and the name of that piece of work has passed into every European language while losing its meaning along the way. It was called the *masterpiece*. Not a work of genius; an examination script. The Nuremberg locksmiths judged a lock; the Parisian joiners a cabinet; and the guild's assembled masters turned it over, ran their thumbs along its joints, and voted on whether the man who made it could be trusted to make things, unsupervised, in their town, under their collective name, for the rest of his life. A licence, in short — but a licence you could hold in your hands. It is worth pausing on how strange and how sane this system was. Seven years watching a master's hands; three years watching a hundred other masters' hands; one object, judged by hands. At no point in

the entire process was anything of importance written down. That was not an oversight.



It was not written down because it could not be. This is the deep fact about skill, and it took until 1966 for anyone to state it in a form that stuck, when a Hungarian-British polymath named Michael Polanyi — a first-rate physical chemist who gave up chemistry for philosophy, to the loud dismay of his colleagues — compressed it into the six words at the head of this chapter: *we can know more than we can tell*. You know how to ride a bicycle. Very well: tell me. Which way do you turn the handlebars when the bicycle begins to fall to the left? (You turn them *left* — into the fall — and if as a child you had done what any written instruction would plausibly have told you, you would have hit the ground. Your body knows this. You, in the telling part of you, almost certainly did not.) You can recognize your mother's face among a million faces; you cannot say how. The blacksmith knows, through the hammer, the moment the iron stops being willing; the baker's hands know when the dough is proved; and none of them can tell you, not because they are inarticulate but because the knowledge does not live in the part of them that talks. Polanyi called it tacit knowledge, and my favourite demonstration of it is the chicken-sexers of Japan. Determining the sex of a day-old chick is commercially urgent — you do not want to feed cockerels that will never lay — and visually it is all but impossible: the relevant cues are subtle beyond description. In the 1920s Japanese researchers worked out that it could be done at all, and a school was soon founded at which students learned to do it the only way it can be learned: the novice guesses, chick after chick, thousands of chicks, while a master stands behind him saying only *yes* or *no*, until one day the novice simply *sees* it, at a glance, at a rate of a thousand chicks an hour with tiny error, and — here is the point — he can no more explain his skill to the next novice than his master could to him.

The knowledge went from hand to hand, so to speak, without ever once passing through words. Now you know what the seven years were for, and the wandering, and the masterpiece. The guild system was not a conspiracy with a craft attached. It was a machine for copying what cannot be written down, and the copying medium was a decade of a human life, and there was no other medium. Knowledge that could be told went into books, and Diderot — you may remember from the last chapter — sent his engravers into the workshops to catch the trades on paper, plate after magnificent plate, and the plates show you everything except the thing itself: how hard to push, when to stop, what *done* feels like.

Hold that thought, and now hear the hinge on which this chapter turns. In the 1980s, when artificial intelligence meant "expert systems," teams of knowledge engineers were sent out, Diderot-fashion, to interview the experts — the physicians, the geologists, the loan officers — and write their expertise down as rules, and feed the rules to the machine. The projects failed, on the whole, and they failed for a reason Polanyi could have given them in advance: the experts did not know what they knew. Ask a physician how she knows this rash is the bad kind and the honest answer, beneath the retrofitted reasoning, is the chicken-sexer's answer: I just see it. For thirty years this was thought to be the permanent wall — economists called it Polanyi's paradox, and it was a great comfort, because it implied the machines would forever be locked out of everything that mattered. Then the wall fell, and notice *how* it fell, because this is the whole story: nobody breached it by finding better words. The new machines are not told rules at all. They watch. Show a network ten million labelled photographs and it learns to recognize a face — not because anyone articulated what a face is, but the way the chick-sexer's apprentice learned: guess, be corrected, adjust, a million times a second, until the seeing happens. The talking machine at the centre of this book learned to write the same way, from everything we ever wrote. Do you see what this means for our subject? The one thing tacit knowledge required — a living

apprentice, standing at the master's shoulder, watching — is the one thing we have now built artificially. The machine does not read the manual. There was never a manual. The machine serves the apprenticeship — every apprenticeship, all of them at once, at a shoulder that never knew it was being watched — and it serves the seven years in an afternoon. Chapter Two warned you that we were, uncomfortably, the drawboy; here is what that meant — every one of us, pulling the cords, unpaid and unasked, above the loom we didn't know we were working.



Before the machine was an apprentice, though, it was a woman with a pencil, and this part of the story deserves telling at length, because we have built our most powerful technology on the graves of the people it is named after, and hardly anyone knows their names.

The word "computer" entered English in the seventeenth century meaning a person who computes, and for the next three hundred years that is all it meant. When the Astronomer Royal, Nevil Maskelyne, launched the *Nautical Almanac* in 1767 — the annual volume of pre-calculated lunar positions that let a ship's officer find his longitude, and incidentally let Britannia rule the waves — the calculations were done by a cottage industry of human computers scattered across England, working at home, paid by the job, posting their sheets of figures back to Greenwich. Maskelyne, no fool, had every table computed twice, by two computers who did not know of each other's existence, and employed a third person, the comparer, to catch the discrepancies — a protocol, you will notice, that we have lately reinvented for checking the arithmetic of machines. Among his steadiest computers for decades was one Mary Edwards of Ludlow, who did the work in her husband's name until the day he died, whereupon it emerged that she had been doing most of it all along; to its credit, the Board of Longitude looked the other way and kept the cheques coming, and she computed on, under her own name, for thirty years more. The

work was piecework, it was respectable, it was poorly paid, and it was *skilled* — a good computer carried an arsenal of tricks, checks and instincts for the plausibility of a number. Gauss, who computed prodigiously himself and enjoyed it, is supposed to have remarked that nothing exposes the lack of mathematical education so surely as meaningless precision in calculation — and whether or not he said it (I pass it on, as usual, at gossip speed), it is the human computer's epitaph in advance: knowing *which* figures matter was precisely the part of the job that was never in the arithmetic.

The trade found its cathedral in, of all places, an observatory. In 1877 Edward Pickering became director of the Harvard College Observatory and inherited a mountain: hundreds of thousands of photographic plates of the night sky, each speckled with stars whose positions and brightnesses had to be measured, logged and classified — a labour beyond any conceivable staff of gentleman astronomers. The story goes that Pickering, exasperated with a male assistant, declared his Scottish maid could do better — and the story is probably polished, as such stories are, but the fact beneath it is solid, because he then went home and hired her. Williamina Fleming had come from Dundee, been abandoned by her husband, and gone into service; within a few years of joining the Observatory she had devised a working classification of stellar spectra, catalogued more than ten thousand stars, and spotted, on a plate in 1888, the curl of darkness we now call the Horsehead Nebula. Pickering hired more women — they were painstaking, they were available, and (he was a man of his time, and said so) they were cheap: twenty-five to thirty-five cents an hour, less than a good secretary — and the group was known around Cambridge, with a levity I decline to launder, as "Pickering's harem." Remember what they actually did. Annie Jump Cannon, profoundly deaf since young womanhood, classified the spectra of well over three hundred thousand stars, at a rate of three stars a minute, and the scheme she distilled — O, B, A, F, G, K, M — is the one every astronomy student on Earth still memorizes. And

Henrietta Swan Leavitt, set to the drudgework of measuring variable stars in the Magellanic Clouds, noticed in the columns of her own figures that the brighter such a star, the slower its pulse — a relationship which turned out to be the first yardstick that could reach beyond the galaxy, the tool with which Hubble, a few years after her death, measured the universe and found it running away from us. A Swedish mathematician began proceedings to nominate her for the Nobel Prize in 1925, only to learn she had died, at fifty-three, four years earlier, having spent nearly her whole career on the books as a computer. When people tell you the routine work is beneath human beings and the machines are welcome to it, it is worth remembering that the expanding universe was discovered *inside* the routine work, by a computer who was paid thirty cents an hour to notice nothing.

The trade's finest hour and its funeral came together, in the 1940s, and the two events were the same event. At Bletchley Park the bombes and the Colossus machines were tended around the clock by hundreds of young women of the Women's Royal Naval Service — Wrens — who kept the secret, most of them, for thirty years afterwards, telling curious grandchildren they had been clerks. In the American war effort, computing rooms full of women ground out the ballistics tables the gunners needed — you watched them at Aberdeen in Chapter Three; on the census forms and the government pay-scales, "computer" was still an occupation, and would remain one into the 1950s — sub-professional grade, note, whatever the mathematics involved. At Langley in Virginia the aeronautics laboratory kept a segregated computing pool, the West Area Computers, staffed by black women mathematicians; one of them, Katherine Johnson, calculated the trajectory for the first American in space, and in February 1962, when the orbital flight of John Glenn depended on a new-fangled electronic machine, an IBM prone to hiccoughs, Glenn declined to fly on its unsupported word. *Get the girl*, he told the engineers — Johnson, he meant — to run the numbers; *if she says*

they're good, then I'm ready to go. She ran the machine's trajectory by hand, through her own equations, and pronounced the numbers good, and Glenn went. Savour that arrangement, because it will come round again before this book is done: the machine computes; the human being, with the craft in her hands and her name on nothing, decides whether to believe it. And then, within a decade or two, the title crossed over. The electronic machines that had been named after the women — "electronic computer" was originally a *distinction*, you understand, like "horseless carriage" — became simply computers, and the women became, in the litany of the age, redundant; though some, and the justice of this is exquisite, walked straight through the fire into the new trade on the other side of it: when the ENIAC was unveiled to the press in 1946, the six people who had actually wrestled it into performing — who had invented, from nothing, on the fly, much of what programming even *means* — were six former human computers; and when the celebrating was done that evening, the men went on to the dinner, and the six women, by every account that survives, were not invited. The first job the computer ever took was the computer's. And the first people who ever taught the machines their trade were the last people whose trade it had been.

★

The switchboard operators went the way Strowger had ordained, though it took the better part of a century — at mid-century the Bell System was still one of the largest employers of women in America, hundreds of thousands of them, wired into the boards like organists, and the last manual exchanges dragged on into living memory — and with them went a skill nobody had ever bothered to classify: the operator was the network's error-correction, its emergency service, its memory of who was at what number, and "number, please" was the sound of a system with a mind in the middle of it. The lamplighters went too, and they did not go quietly: in April 1907 New York's lamplighters struck, and for a few nights

whole districts of the city stood dark — a strike, if you think about it, of men whose entire craft was the nightly refutation of darkness — and the city improvised, and the electric lamps were coming anyway, lamps that needed no ladder, no wick, no man; though I am pleased to report, for the collectors of survivals, that a tiny corps of lamplighters tends London's last fifteen hundred gas lamps to this day, on contract, with ladders, in Westminster, and that nobody who has seen the light they give has ever mistaken it for the electric kind. The typing pool rose and fell inside a single century: the typewriter, arriving in the 1870s, created — almost overnight, as these things go — the first mass respectable indoor occupation for women, so completely that by 1930 the census found typing and stenography to be women's work at upwards of ninety-five per cent, and the word "typewriter" itself, in the early decades, meant the woman as well as the machine, which tells you everything about how the age saw her; then the word processor arrived in the 1970s and the pool drained in twenty years. Each of these is remembered, when it is remembered at all, as a story about employment. Look again, and every one of them is a story about *skill*: a real, difficult, cultivated human capability — the operator's ear, the lamplighter's round, the typist's hundred-word-a-minute hands — becoming worthless not gradually, with time to adjust, but at the speed a purchasing department can sign an order.

And if you want to know what such a skill looks like from the inside while it is still alive and priced accordingly, go to London and watch a young man on a moped, in the rain, with a laminated map clipped to his handlebars, muttering. He is doing the Knowledge, and has been, most likely, for three years. To drive a black cab in London you must pass an examination instituted in 1865 and barely softened since: the Blue Book's three hundred and twenty listed runs, and behind them the twenty-five thousand streets and perhaps a hundred thousand "points" — hotels, hospitals, statues, stage doors — within six miles of Charing Cross, all of it to be carried in the head, none of it on paper. The examinations are one-

man oral ordeals called, with beautiful understatement, "appearances": an examiner across a desk names two points he may well have chosen for their obscurity — a chapel here, a club there — and the candidate must recite the run between them, turn by turn, street by street, *comply this, leave that on your left*, while the examiner listens for the one wrong turning. Most candidates take three to four years. Most who start never finish. It is the last full-dress apprenticeship examination in the Western world that ordinary people ever encounter — the masterpiece, surviving as a recitation — and every element of it, the years, the wandering (on a moped), the master across the desk, the licence at the end, is pure fourteenth-century guildcraft, alive in the age of the satellite navigator that can do the whole thing, for anyone, for nothing. The cabbies know exactly what is happening to the value of what they carry in their heads. Ask one. He will tell you, eloquently, all the way to your destination, by a route the satellite would never have found.



Sometimes the machine does not even need to imitate the skill; it abolishes the terrain the skill grew on. In April 1956 a converted tanker called the *Ideal-X* sailed from Newark, New Jersey, for Houston, carrying fifty-eight aluminium boxes on its deck, and the man who had put them there was not a shipping magnate but a trucker: Malcom McLean, of North Carolina, who had spent his working life waiting at piers watching gangs of longshoremen swing his cargo ashore sack by sack and barrel by barrel, and who had concluded — the perspective of the outsider, once again, being the sharpest tool in this book — that the object of the exercise was not to load ships but to move goods, and that the entire waterfront way of life was, from the goods' point of view, a bottleneck with a culture attached. The culture was a genuine craft culture, mind. A master stevedore's stowage plan was a three-dimensional puzzle balancing weight, sequence of ports, the chemistry of adjacent cargoes and the temper of the

sea, learned in the body over years; dockers could read a sling-load's balance at a glance; and the docks bred the densest working-class communities in the industrial world, with their own words, their own pubs, their own dynasties. The box did not negotiate with any of that. Loading loose cargo in 1956 cost, by the standard reckoning, five dollars and eighty-six cents a ton; loading the same ton in a container cost less than sixteen cents — a collapse of some ninety-seven per cent, deeper even than print's, and the great dock strikes of the sixties and seventies, on both sides of the Atlantic, long, bitter, and doomed, could do no more about it than the scribes could do about Gutenberg — or than the silk-weavers of Lyon could do about Jacquard's cards, who are said to have hurled their wooden shoes, their *sabots*, into his looms and thereby given us the word *sabotage*, a story the etymologists assure me is too good to be true, which has never yet stopped a story. Within a generation the upriver Port of London — for two centuries the busiest thing on Earth — was simply gone, its basins filled or yachted, its trade moved downstream to a crane-forest at Tilbury and Felixstowe where a handful of men in a control room land more tonnage in a day than ten thousand dockers once handled in a month. I said the box abolished the terrain, and I meant it precisely: the stevedore's skill was not out-competed. The situation it applied to was removed from the world.

Yet the ledger of these collapses is not simple, and the honest historian must now produce the century's favourite counter-example and then, reluctantly, complicate it. The automated teller machine arrived in 1967 — the first one at a Barclays branch in Enfield, north London — and everyone could see what it meant: the bank teller, whose defining task was handing over cash, was finished. What followed instead has become the standard economist's parable, and it deserves its fame: over the next four decades, as the machines spread by the hundred thousand, the number of human bank tellers in America *rose*. The machine took over the counting of cash; counting cash had been most of the cost of a branch; branches

became cheap; so the banks opened more branches, which needed tellers — fewer per branch, more overall — and the teller's job quietly transformed from counting into talking: advice, sales, the ironing-out of trouble, the parts of banking that walk on two legs. Automation, runs the moral, does not eliminate work; it relocates it into whatever the machine cannot do, and the total, historically, has gone up. It is a true story and a good moral and I have dined out on versions of it myself. Now the complication, which the parable's retailers usually omit: it stopped working. Around 2010 the number of tellers peaked, and it has fallen every year since, because the telephone in your pocket removed the reason to visit the branch at all, and the branches themselves — the terrain — began to close. The parable, in other words, held for exactly as long as the machine could do only *part* of the job. That is the general rule hiding inside the reassuring anecdote, and it is worth stating plainly, because the whole comfort of the last forty years of economics rests on it: partial automation multiplies the human remainder; the trouble starts when the remainder itself is next on the list. Every profession currently consoling itself with the ATM story should ask which half of the parable it is living in, and check the date.



There was, historically, another road, and it is time to give a proper entrance to the man who built it — you glimpsed him in the last chapter, a young radar technician reading Vannevar Bush in a Red Cross library — because the choice between the two roads is about to be put to us again, harder than ever before.

In December 1968, at a computer conference in San Francisco, a soft-spoken engineer from the Stanford Research Institute named Douglas Engelbart sat down on stage beneath a twenty-foot screen, settled a microphone over his head like a fighter pilot, and for ninety minutes showed a thousand of the world's best computer people things none of them had ever seen: text edited freely on a screen; documents linked to

other documents; two people, thirty miles apart, working on the same file at once, their faces inset on the display; and, steering it all, under Engelbart's right hand, a small wooden box on wheels that his team called — nobody now remembers exactly who first said it — a mouse. The audience stood and cheered. The performance has been known ever since as the Mother of All Demos, and virtually everything in it sits on your desk today. But the demo was the least of it. Engelbart was a man with exactly one idea, conceived in 1950 when as a young engineer, newly engaged and appalled to find he had no goals worthy of a lifetime, he reasoned his way in a matter of months to this: the world's problems were compounding faster than human institutions could cope, so the highest-leverage thing any engineer could do was raise the capability of the human mind itself — *augment* it, with machinery, the way the lever augments the arm. Not replace. Augment. He wrote the manifesto in 1962, built the lab, built the demo — and note that at the very same moment, a few miles away, the other road was being surveyed: the artificial-intelligence laboratories of the day took it more or less as read that the machine's destiny was to do the thinking itself, and the two camps regarded each other across the Bay with polite incomprehension. History, with its usual taste for irony, then arranged matters so that Engelbart's gadgets conquered the world while his idea was forgotten — SRI licensed the mouse to Apple for a reported forty thousand dollars, flat, against the billion-odd mice the world went on to buy, and Engelbart, who never saw a royalty, spent his later decades as a revered, unfunded prophet — but the idea is the thing, so let us give it its proper test. Which road wins? For a while, the answer looked wonderfully clear, and it came from chess.

You will recall from Chapter Five the evening in 1997 when Kasparov lost to Deep Blue and the world declared thinking abolished. What the world missed was what Kasparov did next. Within a year, in León in Spain, he staged a match of what he called advanced chess: each player at the board with a computer at his elbow — man *and* machine against man and

machine, the centaur configuration, Engelbart's idea in sixty-four squares. And when the format spread to open tournaments, something glorious happened. In a famous freestyle event in 2005, teams of grandmasters with supercomputers were beaten — soundly — by a pair of young American amateurs, rated nowhere, running three ordinary personal computers, whose secret was simply that they knew *their machines* better: when to trust the silicon, when to overrule it, how to arbitrate when the engines disagreed. Kasparov drew from this a law that deserves to be carved somewhere: weak human plus machine plus better process beats strong human plus machine plus worse process — and, for a golden decade or so, beats machine alone. The centaur was the strongest player on Earth. Augmentation had beaten automation in a fair fight, and those of us of an optimistic disposition (I own a television historian's supply) took comfort: there, surely, was the future of every profession — the doctor-centaur, the lawyer-centaur, judgment enthroned above computation. Well. I am obliged to report what happened next, which is that the engines kept improving and the humans, being human, did not, and the gap between what the machine suggested and what the best human could add narrowed, and narrowed, until — by somewhere in the middle of the 2010s, and beyond argument after the self-taught engines of 2017 — the human half of the centaur had become, at chess, a net passenger: his interventions now subtract. The centaur was not a destination. It was an *interval* — the stretch of years between the machine becoming useful and the machine no longer needing the rider — and at chess the interval lasted roughly fifteen years. The question that should be keeping the professions up at night is not whether they get a centaur phase. They will; we are in it; it is lovely; I am living in it as I write this sentence, machine at my elbow like a good second. The question is the length of the interval, and the only honest answer the historian can give is: it varies, nobody knows, and at chess — a narrow, closed, scorable world, everything your profession is not, you hope — it was fifteen years.



Which brings us, with a certain grim relish, to the toolmakers themselves.

There is a special flavour to what is now happening to computer programmers, and to savour it you need one fact about their history: programming is the only craft in this book that has spent its entire existence automating *itself*. The ENIAC women programmed with cables and switches; then assemblers automated the cables; then Grace Hopper — the American naval officer with a pirate's disregard for the impossible — built the first compilers in the early 1950s, so that a human could write something like English and the machine would write the machine-code, and she was told, she cheerfully recorded, that this was ridiculous, since computers could only do arithmetic. Every rung of the ladder — compiler, high-level language, database, spreadsheet — automated away the programmers below it, and every single time the result was more programmers, not fewer, because each rung made computing cheap enough for new problems, in the best ATM-parable style. The craft grew fat on the pattern and made a doctrine of it. And then, within the last few years, the pattern arrived at its fixed point: the talking machine writes the code itself, from a conversation, from a wish — Hopper's dream completed and inverted in a single gesture — and the profession that automated everyone else has discovered recursion in the one place it never wanted to find it. The junior programmer now watches a machine do, in forty seconds and rather well, the work that was to have been his apprenticeship; his senior watches the machine do a growing share of hers, and reflects — she is trained to notice such things — that she has seen this loop before and knows how it terminates. The illustrators are ahead of them in the queue, and their case is the purest in the whole chapter, because what the image machines did is precisely, structurally, what every apprentice since Lascaux has done — learn by looking at the masters' work, millions of instances, until the style lives in the hand — done at a scale and speed no

guild ever imagined, and without the one thing every guild demanded in exchange: the indenture. The bond, the fee, the seven years' service, the sworn secrecy — the whole apparatus existed to ensure that a master's manner of working could not simply be *taken* by anyone with eyes. When the machine apprenticed itself to every living illustrator at once, uninvited, and then set up shop next door selling their styles for pennies, it committed no novel sin. It committed the oldest sin in the craft economy, the exact one the guilds were invented to police — and found the guilds long abolished and the police not yet hired. And behind the illustrators, the paralegals and the junior lawyers, whose mountains of discovery documents — the thousand-box warehouses of subpoenaed paper through which the young were traditionally marched for their formative years — are exactly the kind of reading the machine does best, at midnight, without overtime, without ever once becoming a partner.

Notice what all three cases have in common, because it is the sharpest point this chapter has to make, and I have been circling it long enough. In every one of them, what the machine does first and best is *the junior work*. Of course it does: the junior work is by definition the codified, repeatable, watchable part of the craft — the part you can learn from examples, which is the machine's whole method of being. The senior work — judgment, taste, the sense of when the standard answer is wrong — remains, for now, human, and every profession is busy reassuring itself with exactly that observation. But stay with the thought one step longer than is comfortable. Where do seniors come from? They come from juniors. And what turns a junior into a senior is — every master in this chapter has already told you — not the lecture, not the manual, but the years of supervised tedium: the document review, the routine radiographs, the small bug fixes, the thousandth chick. The tedium was never a regrettable inefficiency in the training. *The tedium was the training* — the modern remnant of the seven years and the wandering, the only medium through which the untellable part of the craft has ever passed from one nervous system to

another. So when the machine takes the junior work, it does not merely displace this year's juniors. It quietly dismantles the factory that makes seniors, while the current stock of seniors — trained the old way, at full price, before the flood — carries on so smoothly that nothing seems wrong for a decade or two. Then they retire. I have no cure to sell you; nobody has one yet, and I distrust everyone who claims otherwise; but I can tell you the problem has been rehearsed, in miniature, in the most instrumented workplace on Earth. Aviation automated the flying of aeroplanes so successfully that airline pilots now hand-fly a few minutes per flight, and as long ago as 1983 an engineer named Lisanne Bainbridge laid out, in a short paper with the deadpan title "Ironies of Automation," the trap in full: the better the automation, the rarer and more atrophied the human skill, and therefore the worse the human performs on precisely the occasions when the automation gives up and hands back control — which are, of course, the worst occasions. Airline training films took to warning pilots against becoming "children of the magenta line," docile followers of the computed track; and one June night in 2009, over the mid-Atlantic, an airliner's speed sensors iced over, the autopilot disengaged as designed, and three qualified pilots, presented with a recoverable situation and several minutes in which to recover it, held the aeroplane in a stall all the way down to the sea. The accident report is careful and technical and says many things, but the thing it says to *this* chapter is simple: the juniors had stopped doing the junior work long before anyone worked out where the seniors would come from. Two hundred and twenty-eight people paid for the rehearsal. The play proper — medicine, law, engineering, code — is now in production, everywhere at once.

And when a chain of hands breaks, it breaks for good; this too has been rehearsed. Skills that took centuries to accumulate go extinct, quietly, in a single unrenowned generation — the working sailmakers who could cut a full suit of square-rig canvas by eye, down now to a scattering of tall-ships specialists; the wheelwrights and the master coopers, of whom Britain

retains a number you could seat at one table; hand typesetting, five hundred years old, which died in the composing rooms within a decade of the computer's arrival — at the *New York Times* the last hot-metal edition went to press on the second of July, 1978, and they made a little film that night of the Linotype men, some of them third-generation, setting their final slugs, and the film takes its title from the two nonsense words a Linotype man ran his finger down the keyboard to cast when a line was spoiled, words which had appeared by accident in the world's newspapers for a century and would now never appear again: *etaoin shrdlu*. There is today an organization in Britain that publishes, every two years, a Red List of endangered crafts, in frank imitation of the one we keep for birds — currently well over a hundred entries, from arrowsmithing to withy-work, each with its census of remaining masters, some of those censuses reading *one*. We have, in other words, begun keeping conservation records for human skill. I leave it to the reader to decide whether that is reassuring.



So we arrive at the question the whole chapter has been walking towards, and it is not the question you were promised by the newspapers. Not: what will the machine learn to do? The answer to that one is now clear enough — nearly anything it can watch us do in sufficient quantity, the tellable and the untellable alike. The live question is the mirror of it: what will we *choose to keep doing ourselves*, when we no longer strictly need to — and how does a civilization even go about making a choice like that? It sounds impossibly novel. It is not. There is a community in North America that has been making exactly this choice, deliberately, tool by tool, for a century, and it is the community every technologist was raised to laugh at.

The Amish do not, whatever the cartoons say, reject technology; a modern Amish workshop is full of pneumatic tools, and an Amish farmer may hire a car and driver on Tuesday whose twin he declined to own on

Monday. What the Amish actually do is *negotiate* with technology, and the negotiation has rules. When a new tool arrives, they do not ask the axemaker's question — what can it do? — they ask a different one: what will it do *to us*? To the family table, to the visiting of neighbours, to the young; and they run trials, and they watch, and then the community decides, and the decision differs from district to district because the communities differ. The telephone is the famous case: they saw, a century before the smartphone made it obvious to the rest of us, that a telephone in the kitchen means the end of the walk to your neighbour's kitchen — so they put the telephone in a shanty at the end of the lane. Not banned: *placed*. The tool admitted, the interruption refused. You may find the theology quaint; their sociology has just spent a hundred years being proven right, and I confess that a man who has written nine chapters about tools remaking their users without so much as a by-your-leave finds it hard to condescend to the one society on record that reads the small print first, out loud, in committee.

The rest of us will not adopt the horse and buggy, and need not. But the market — that blunt, honest instrument — suggests the rest of us have already begun, in our disorganized way, to make Amish-shaped choices, because a curious thing happens to the handmade whenever a machine perfects the making. In the 1970s the quartz movement arrived and did to the Swiss mechanical watch what the container did to the docker: a quartz watch was more accurate, more reliable and a fiftieth of the price, and Swiss horology collapsed, shedding two workers in every three. And then — no economist predicted it — the corpse sat up. The mechanical watch was reborn, not as a timekeeper but as the *proof of hours*: what the buyer of a hand-finished movement is purchasing, knowingly, is precisely the fact that a human being spent human time on it, the one ingredient the machine's version, by definition, cannot contain; and the Swiss industry now earns more from its obsolete technology than it ever earned when the technology was necessary. The gramophone record, flattened by the

compact disc and then by the stream, returned the same way — by the early 2020s vinyl was outselling the CD again for the first time since the 1980s, bought largely by people with unlimited free access to the same music, which tells you it is not the music being bought. The craft brewers, the sourdough bakers, the letterpress printers of wedding invitations at twenty times the laser price: everywhere the machine's version becomes perfect and free, a market grows up beside it, small but insistent, for the version with a person still in it — imperfect on purpose, expensive *because* expensive, the flaw functioning as the proof, the way the torn edge of the indenture was the proof. That market will not employ everyone; let us not sentimentalize it; the sailmakers' Red List entry is not rescued by artisan gin. But it demonstrates the thing this chapter most needed demonstrating: that "the machine can do it" and "we shall therefore stop doing it" are separate sentences, and the second one is ours to write, craft by craft, the way the Amish write it — deliberately, or the way the rest of us do — by default, and too late, and then expensively backwards.

Because here is what we would be choosing to keep, and it is more than nostalgia. A craft practised is a way of knowing the world through the hands, and we have Polanyi's word — and the chicken-sexers', and the cabbies', and Katherine Johnson's — that this knowing is real knowledge, unduplicated anywhere in the telling kind. And a craft *transmitted* is something more again. Go back, one last time, to the beginning of this book: to the first chapter's first scene, the old hands guiding the young hands on the flint, the precise sequence of blows that some believe demanded — and thereby created — precise, sequential speech. If that account is even half right, then master-and-apprentice is not merely the oldest institution we have. It is the institution that *made* us; language itself was its first graduation piece; every other structure in this book — the guild, the university, the residency, the Knowledge, the observatory full of underpaid women teaching each other the stars — is that first workshop, refounded. For a million years the chain of hands has never once been

broken. Now we have forged a new link of a new metal, an apprentice of unbounded diligence that learned the untellable by watching all of us at once, and it has submitted its masterpiece to the guild, and the assembled masters — that is you and I — turn the thing over, and run our thumbs along the joints, and find the work disturbingly good. And here the old rules of the workshop fail us, because the masterpiece this apprentice offers is *imitation itself* — work in any master's manner, any hand, any voice, mine included, indistinguishable from the original to any thumb in the room. The guilds' whole world rested on being able to tell whose work was whose; the indenture authenticated itself by a torn edge no forger could match. What happens when nothing does — when the copy of the hand is perfect, and the voice on the recording may be no one, and the torn edges of the world all fit? That is not a question about work at all. It is a question about truth, and it has been waiting for us, patient as an examiner, in the next chapter.

CHAPTER TEN

True Enough

*Falsehood flies, and the Truth comes limping after it; so that when
Men come to be undeceiv'd, it is too late; the Jest is over, and the
Tale has had its Effect.*

— JONATHAN SWIFT, *THE EXAMINER*, 1710

In the summer of 1917, in the village of Cottingley near Bradford, a nine-year-old girl named Frances Griffiths kept coming home with wet shoes and a wet pinafore, having fallen — again — into the beck at the bottom of the garden. Asked why she would insist on playing down there, she gave the excuse that has launched a thousand parental sighs: she went to see the fairies. Her sixteen-year-old cousin Elsie Wright backed her up, and the two of them, tired of being laughed at, borrowed Elsie's father's camera — a Midg quarter-plate, one exposure loaded — and disappeared down to the beck for half an hour. When Arthur Wright developed the plate in his cellar darkroom that evening, there was Frances, chin on hand,

gazing coolly at the lens while four winged figures danced on the foliage in front of her. Mr Wright, an electrician and a sensible man, took one look and said the figures looked like bits of paper. Which, of course, they were. Elsie had copied them from a picture book, given them wings, cut them out, and fixed them upright with hatpins. Her father searched the girls' bedroom and the banks of the beck for scraps and found nothing, because the fairies had gone down the stream, and there the matter should have died: a family joke, two girls giggling behind their hands, one wasted photographic plate.

Instead, three years later, the creator of Sherlock Holmes told the world it was true.

The chain of events by which that happened is worth following, because it is the whole of this chapter in miniature. Elsie's mother Polly, unlike her husband, believed. In 1919 she went to a Theosophical Society lecture in Bradford — the subject, as luck would have it, was fairy life — and mentioned the photographs. Prints climbed the theosophical ladder until they reached a leading light of the movement named Edward Gardner, who sent them to a photographic expert named Harold Snelling, who pronounced them genuine, single-exposure, untampered-with photographs of whatever was in front of the camera — which was perfectly true, and perfectly useless, since what was in front of the camera was paper. And then the prints reached Sir Arthur Conan Doyle, who had been commissioned by the *Strand Magazine* to write an article about fairies for the Christmas number and could scarcely believe his luck. Doyle did his due diligence, after a fashion: he showed the pictures to the physicist Sir Oliver Lodge, who smelt a hoax, and sent them to Kodak, whose examiners found no evidence of darkroom trickery but declined, with the weary tact of men who knew exactly what was being asked of them, to issue a certificate that fairies were real. Doyle published anyway. The *Strand* of December 1920 — "Fairies Photographed: An Epoch-Making Event" — sold out in days.

Now, why? How does the inventor of the most relentlessly rational detective in literature, the man who taught the whole world to say "when you have eliminated the impossible, whatever remains, however improbable, must be the truth," come to stake his reputation on hatpins and paper? Because of what had happened to him in the meantime. Doyle's son Kingsley had been wounded on the Somme and died of pneumonia in 1918, part of the great harvest of sons that made the séance parlours of the 1920s the fullest rooms in Europe; Doyle had thrown himself into spiritualism with all the formidable energy he had once given to Holmes, and by 1920 he needed the fairies to be real, because if the fairies were real then the walls of the material world had a door in them, and Kingsley was on the other side of it. (There is a further turn here, almost too neat to bear: Doyle's own father, Charles Altamont Doyle, had spent his last years in Scottish asylums filling sketchbooks with paintings of fairies. And the picture book from which Elsie copied her dancing figures — *Princess Mary's Gift Book*, 1914 — contained, a few pages from the illustration she traced, a story by one Arthur Conan Doyle. The great man was deceived by paper cutouts sourced, in a manner of speaking, from his own bookshelf.)

The two cousins kept the secret for over sixty years, mostly, they said later, out of embarrassment — not their own, but everyone else's. "Two village kids and a brilliant man like Conan Doyle — well, we could only keep quiet," said Elsie, when the pair of them finally confessed in the 1980s, two respectable grandmothers explaining the hatpins to the television cameras. And here is the detail I want you to keep in your pocket for the rest of this chapter: Frances, who admitted freely that four of the five photographs were faked, insisted to her dying day that the fifth one was real.

The lesson of Cottingley is not that photographs can be faked. Everybody in 1920 knew photographs could be faked; the photographic press said so, loudly, at the time, and the girls' own father had said it in the first five minutes. The lesson is that the fake does not need to be good. It

needs to be *wanted*. The doctored image does not seek out the sharpest eye; it seeks out the believing one, and when it finds it, no amount of Kodak will help. Every forgery in this chapter — and this chapter is a rogues' gallery, seven centuries of them — succeeded not because the forger was brilliant but because the audience was ready. Keep that in mind as we go, because we have just built a machine that can manufacture the wanted thing, on demand, for everyone, at once — evidence woven to order, in anyone's voice, wearing anyone's face. Including, as we shall see, mine.



Long before anyone could doctor an image, they doctored the document, and the greatest doctored document in history propped up an empire for six hundred years before a bad-tempered Italian took it apart with a grammar book.

Some time in the eighth century — nobody knows quite when, or quite where, though the smart money is on the papal chancery or somewhere very cosy with it — somebody composed a document called the Donation of Constantine. It told a wonderful story. The Emperor Constantine, stricken with leprosy, is cured through baptism by Pope Sylvester I, and in his gratitude does the decent thing: he hands the Pope his palace, the city of Rome, all Italy, and indeed the whole of the Western Empire, and moves himself off to Byzantium so as not to crowd the successor of Saint Peter. He even offers Sylvester his imperial diadem, which the Pope modestly declines. For the next several centuries, whenever a pope needed to remind an emperor or a king who really held the deeds to Europe, out came the Donation. It was cited in disputes over Sicily and the Papal States; it was painted on chapel walls; it was, in the lawyer's sense and every other sense, title.

There were doubters, of course. In the year 1001 the young emperor Otto III issued a diploma flatly calling the Donation a fabrication, and even naming the forger — a certain Deacon John, known, pleasingly, as

"John of the mutilated fingers" — but an emperor calling a papal deed forged is a litigant disputing the other side's evidence, and everyone shrugged. The demolition, when it came, came from a new direction entirely. In 1440 Lorenzo Valla — priest, scholar, and possessor of the finest and most weaponized Latin in Europe — was employed at the court of Alfonso of Aragon, who happened to be at war with the Pope over the crown of Naples, so let us not pretend the fact-checker had no client. But what Valla did with the brief was genuinely new. He read the Donation the way a geologist reads a cliff face. And he showed, clause by clause, that the thing could not possibly have been written in the fourth century, because it was full of the eighth. The text has Constantine handing out "satraps" — a Persian office; Rome never had satraps, but the eighth century liked the sound of them. It refers to Constantinople — in a document supposedly executed before the city of that name existed. It describes the imperial diadem as a cloth crown when any fourth-century official would have known it was silk and pearls, it garbles the titles of Roman officialdom, and — the cut that must have hurt most in the chancery — it is written in a Latin so shabby, so anachronistic, so *provincial*, that Valla professionally refuses to believe any servant of Constantine could have produced it. Language, it turns out, is sedimentary. Every document carries the fingerprint of its moment in its vocabulary and its grammar, and no forger can fully scrub it, because no forger knows which of his own words haven't been invented yet.

Valla's little book is the founding act of textual forensics — philology as fact-checking, the first time in European history that a document was convicted by internal evidence, by *style*, rather than by a witness or a confession. Hold on to that idea too, because at the end of this chapter we are going to need it badly: the fake betrayed by the statistics of its own language. And savour the two ironies that follow. First: the Renaissance papacy, having read Valla's demolition of its most treasured deed, responded by hiring him — Nicholas V made him a papal secretary, on the

eminently sensible grounds that a weapon like that is better inside the tent. Second: Valla's treatise circulated quietly in manuscript for decades until a German knight named Ulrich von Hutten had it printed — in 1517, of all the years available — and a copy reached Martin Luther in 1520, and the discovery that the central land-deed of papal power was a forgery landed on that already-smouldering man like a bucket of pitch. The fact-check, seventy years old, detonated the moment print got hold of it. Which tells you something we shall meet again and again: the truth business and the falsehood business use the same delivery vans.

And how, you might ask, did anyone authenticate anything during the long centuries when documents could be conjured like this? The answer is that civilization had already learned the trick it would keep relearning: when the thing itself can be faked, attach it to something guarded. The seal — wax pressed with a matrix so jealously kept that counterfeiting the king's Great Seal was defined in the English Treason Act of 1351 as high treason, dying-badly territory, the same class of crime as levying war on the king. The notary — a profession the medieval world inherited from the Roman *tabellio* — whose elaborate personal sign at the foot of a deed said, in effect, *a licensed and traceable human being stakes his livelihood that he watched this happen*. The signature. And my favourite, because the name gives the whole game away: in 1300 Edward I decreed that no silver was to be sold until it had been assayed and struck with the leopard's head, and the striking came to be done at the Goldsmiths' Hall in London, and the stamp became known as the hall's mark. The hallmark. Notice what a hallmark actually is. It is not a property of the silver; any competent villain can fake a stamp. It is a pointer to an *institution* — the assay office, with its scales and its acids and its records and, by the eighteenth century, the death penalty for counterfeiting the mark. Trust never lived in the metal. It lived in the Hall. Remember the Hall.

Then came print, and the price of the written word collapsed — the economics of that collapse belong to the next chapter, so I'll resist — and with it collapsed something nobody had thought to price separately: the cost of lying to strangers at scale. A manuscript lie travelled at the speed of a scribe. A printed lie travelled at the speed of a thousand sheets a day per press, and the presses did not check references. Within a generation of Gutenberg, Europe was hip-deep in pamphlets, and the pamphlet wars of the Reformation established the enduring rule of the new information order: the fastest, crudest, most infuriating version of events wins the first week, and the first week is when minds are made up. Swift's line at the head of this chapter was not a philosopher's abstraction; it was a working journalist's report from the trenches of 1710.

If you want to know what the new order could do at full throttle, London in the autumn of 1678 will oblige, and it will not be pretty. A failed clergyman named Titus Oates — expelled from school, sent down from Cambridge without a degree, dismissed from a naval chaplaincy in disgrace, and lately in possession of a doctorate from the University of Salamanca that the University of Salamanca had somehow neglected to award him — came forward with a sworn narrative, forty-three articles long and growing, of a vast Jesuit conspiracy to assassinate King Charles II and put his Catholic brother on the throne. It was invention from end to end, cooked up with a fanatical associate named Israel Tonge, and it might have gone nowhere but for a piece of gruesome luck: the magistrate who had taken Oates's deposition, Sir Edmund Berry Godfrey, was found dead in a ditch on Primrose Hill, strangled and run through with his own sword. Nobody has ever discovered who killed Godfrey — the murder remains open to this day — but London *knew*, the way a crowd knows, and the presses told them what they knew, and the country went mad for three years. The Licensing Act lapsed in the middle of the crisis and the unlicensed presses poured out narratives, last dying speeches, revelations, each pamphlet ratcheting the terror a notch higher. Oates was lodged in

apartments in Whitehall with an allowance of twelve hundred pounds a year and hailed as the Saviour of the Nation; he strolled the town with an escort, pointing out traitors, and men died on his pointing. At least fifteen innocent men were executed — the last of them Oliver Plunkett, Archbishop of Armagh, hanged, drawn and quartered at Tyburn in 1681, and canonized as a saint three centuries later, which was handsome but tardy. That is what I mean when I say fake news had a body count long before anybody coined the phrase. When the fever broke and the perjuries unravelled, the state's revenge on Oates was almost as instructive as the plot: convicted of perjury in 1685, he was defrocked, pilloried, whipped through the streets from Aldgate to Newgate and, two days later — the sentence was calibrated to kill him, and somehow didn't — from Newgate to Tyburn, and locked up for life with an annual appointment at the pillory to remind everyone. And then, four years on, the regime changed, and the new one found his story politically convenient after all, and Titus Oates walked out of prison with a pardon and a pension. If you want a one-sentence history of misinformation, there it is: the liar was whipped for the lie, and paid for it too — by different governments, in that order.

But here is the thing about the seventeenth century, and it is the hinge of this whole chapter: the same presses, the same coffee-houses, the same credulous, pamphlet-maddened London, produced — in exactly the same decades — the most successful truth-machine ever built. In 1660 a group of natural philosophers founded the Royal Society and took as their motto *Nullius in verba*: take nobody's word for it. Not the ancients' word, not a gentleman's word, nobody's. Claims were to be demonstrated before witnesses — the Society kept a register, and the experiments were performed in the room — and from 1665 its secretary, Henry Oldenburg, published the *Philosophical Transactions*, in which a claim was dated (so priority could be proved), signed (so credit and blame had an address), and exposed to every rival in Europe with an incentive to demolish it. What grew from Oldenburg's editorial habits eventually hardened into peer

review, but the deep invention was already there in the motto: the Society did not solve the problem of lying by finding trustworthy men. It solved it by building a *procedure that assumed everybody might be wrong*, and made claims fight for their lives in public. And around the wider print trade, more slowly, the same scar tissue grew: the byline, which pins words to a person; the dateline; the masthead with an address you could sue; the libel courts; and eventually that most touching of institutional confessions, the correction column — a standing admission that the paper errs, which, by the marvellous jujitsu of credibility, makes you trust it more. None of this is technology. You cannot point to the machine that is peer review. Print created the truth problem, and what fixed it — to the extent it was fixed, which was never entirely, see Oates, above — was institutions: procedure, accountability, reputation with something to lose. The Hall again, rebuilt in paper.



Photography arrived in 1839 trailing a slogan it never earned: the camera cannot lie. Light itself would now do the testifying — no scribe, no engraver, no fallible human hand between the world and the record. Nature would draw her own picture, and nature was not for hire. It is one of the most beautiful ideas the nineteenth century had, and it lasted, at the outside, eighteen months.

Consider Hippolyte Bayard, a French civil servant who invented, more or less single-handedly and more or less simultaneously with Daguerre, a process for making direct positive photographic prints on paper, and who in June 1839 mounted what was very likely the world's first public photographic exhibition. Bayard had the misfortune to be politically inconvenient: the great astronomer-politician Arago was busy securing a state pension for Daguerre, whose process was to be presented as the official French miracle, and Arago persuaded Bayard to postpone announcing his own — just for a little while, you understand — with the result that

Daguerre got the pension, the fame and the eponym, and Bayard got six hundred francs and a pat on the head. His response, in October 1840, was the most eloquent sulk in the history of art. He photographed himself slumped half-naked against a wall, eyes shut, hands composed — a drowned corpse, fresh from the morgue — and wrote on the back of the print a mock notice explaining that the body was that of Monsieur Bayard, inventor of the process you see before you, whom the government had rewarded with nothing, whereupon the poor wretch had drowned himself; note, the caption adds helpfully, that his face and hands have already begun to decay. There it is. The first staged photograph — the medium caught lying at the very moment of its birth, and lying, please observe, in a good cause and in plain sight: a protest, a joke, a fiction wearing the full uniform of evidence. Photography did not lose its innocence. It never had any. Only its audience did.

And the audience took decades to catch on, because the doctoring professionalized almost immediately. By the 1850s the retouching desk was standard furniture in the portrait studio — nobody's complexion in a Victorian photograph is to be trusted, which is the gentlest possible corruption but corruption all the same, the pixel bought and paid for. By 1857 the Swede Oscar Rejlander had combined some thirty negatives into one vast moral tableau, *The Two Ways of Life*, so seamlessly that Queen Victoria bought a print for Albert; a year later Henry Peach Robinson stitched five negatives into *Fading Away*, a consumptive girl dying between her grieving relations — every figure photographed separately, the deathbed assembled like a jigsaw — and the public was furious, not because the picture was a composite but because it *might not be*, and had made them feel like intruders. And then the fraud industry proper arrived, and its first tycoon was a Boston engraver named William Mumler, who discovered in 1861 — by accident, he always claimed, which is what they always claim — that a carelessly recycled plate would print a faint previous image behind a new portrait. A ghost. Mumler set up shop selling

photographs of the bereaved with their dead hovering behind them, and business, in an America about to lose three-quarters of a million sons, was very good indeed. Note the mechanism — it is Cottingley's, forty years early: the fake finds the believing eye, and the believing eye in 1865 was any eye that had read a casualty list. Mumler's masterpiece came late in his career, around 1872, when a veiled widow calling herself Mrs Lindall sat for him, and the developed plate showed the tall, unmistakable figure of Abraham Lincoln standing behind his wife with his hands on her shoulders. Mary Todd Lincoln — the most famous widow in America, the widow of the most photographed man in America — went to her grave comforted by a double exposure.

Mumler had already been through the courts by then, and his trial repays attention, because it is the moment the law worked out what a photograph actually is. Prosecuted for fraud in New York in 1869, he faced a parade of experts who demonstrated half a dozen ways of making spirit photographs, and a star witness for the prosecution in the shape of P. T. Barnum — Barnum, of all people, testifying against humbug; the century did have a sense of humour — who had commissioned a photographer to fake up a portrait of Barnum with Lincoln's ghost, and entered it in evidence to show the jury how easily it was done. And Mumler was acquitted. Acquitted because showing that a thing *could* be faked is not proof that it *was*, and nobody could demonstrate precisely how Mumler did it. The verdict scandalized the rationalists, but the deeper doctrine that came out of that era was sound, and we are about to need it again, so mark it well: the courts decided that a photograph, on its own, is nothing. It enters evidence only as what the lawyers came to call pictorial testimony — an illustration of what some human witness, under oath, with a name and a face and a punishable body, swears is a fair representation. The photograph never testifies. Somebody testifies, holding a photograph. Evidence, in other words, was never really the artefact; it was always the chain of accountable people attached to the artefact. We forgot that, in the

long comfortable century when faking was expensive. The forgetting was the anomaly, not the rule.

The twentieth century then demonstrated what happens when the doctoring is done not by hucksters but by the state. In the Soviet Union, photographic retouching became a permanent branch of government. Trotsky, who stood beside Lenin in half the iconic photographs of the Revolution, was scratched, painted and cropped out of them one by one as he fell from favour, until whole speeches were being delivered to crowds from which history's second-most-important listener had simply evaporated. Nikolai Yezhov, the murderous little chief of the secret police — they called him the bloody dwarf, and he earned it — strolls along the Moscow–Volga Canal at Stalin's elbow in the official photograph of 1937; after his own turn against the wall in 1940, the photograph was reissued and Yezhov had become water and railing. The retouchers, it must be said, were craftsmen of a high order; the surviving before-and-after pairs, collected by the historian David King under the perfect title *The Commissar Vanishes*, are seamless. And the erasure did not stop at photographs. When Lavrenty Beria was shot in 1953, subscribers to the Great Soviet Encyclopedia — so the story goes, and it sounds too good to be true, except that the replacement pages survive in libraries for anyone to check — received instructions to remove the article on Beria with a razor blade and paste in its place an unexpectedly generous new article on the Bering Sea. The state was not merely lying about the present. It was issuing recall notices on the past — and it teaches us the limit case: where one institution controls the entire chain, from camera to archive to encyclopedia, evidence dies, however honest the emulsion. The Hall guards the mark; but who guards the Hall?

Democracies, before we congratulate ourselves, ran the same play at least once with the volume turned down. Four days before the British general election of October 1924, the *Daily Mail* published a letter, supposedly from Grigory Zinoviev, head of the Communist International in

Moscow, to the British Communist Party, urging sedition in the armed forces and cheering on the Labour government's treaty with Russia — under the tender headline "Civil War Plot by Socialists' Masters". Labour duly lost, the Conservatives romped home, and the letter was a forgery — cooked up, on the balance of the evidence assembled by the Foreign Office's own historian when the files were finally opened in the 1990s, by White Russian émigrés, and passed with suspicious smoothness through the hands of the intelligence services and into the press at the precise moment of maximum damage. Did it swing the election? Honest scholars doubt it — Labour's own vote actually rose; it was the Liberal collapse that did the arithmetic — but that is almost beside the point. What the Zinoviev letter poisoned was deeper than one election: it taught half of British politics, for two generations, that the secret state and the press could not be trusted not to conspire against it, and that lesson, once learned, does not care whether the next document is genuine. File that thought under "liar's dividend"; we shall come back to it.



On the evening of Sunday, 30 October 1938 — the night before Hallowe'en — the Columbia Broadcasting System carried, in its regular Mercury Theatre slot, an adaptation of H. G. Wells's *The War of the Worlds*, produced and narrated by a twenty-three-year-old named Orson Welles. What made it art was the form. The first half of the programme was constructed not as a play but as an evening of broadcasting: dance music from a hotel ballroom, interrupted by a news bulletin, back to the music, interrupted again — a weather notice, a Professor at Princeton, a remote pickup from a farm at Grover's Mill, New Jersey, where something has landed — and the reporter on the scene describes the cylinder unscrewing, and the thing that comes out, in a voice that comes apart exactly the way a real voice had come apart eighteen months earlier when the *Hindenburg* burned at Lakehurst; the actor, Frank Readick, had spent

the afternoon playing the recording of Herbert Morrison's "oh, the humanity" over and over, learning what genuine horror does to a man's breathing. Welles had discovered something for which we still don't have a comfortable name: you do not need to forge the *content* of truth if you can forge its *format*. The bulletin, the interruption, the correspondent, the dead air — by 1938 those signals meant "this is really happening" as firmly as wax and parchment had once meant a king's will. The Martians were fiction; the *grammar* was counterfeit; and some listeners, tuning in late and missing the announcements, took the grammar at its word.

The next morning's newspapers reported a nation in blind panic — stampedes, heart attacks, families fleeing into the night with wet towels over their faces, a tidal wave of terror from coast to coast. It is one of the great set-piece stories of media history, and I have told it myself, with relish, more than once. So it is with a certain professional wince that I pass on what the scholarship of recent decades has established: the panic, to any respectable approximation, did not happen. The audience ratings taken that very night suggest that only a tiny fraction of American radios were tuned to the Mercury Theatre at all — most of the country was listening, sensibly, to the ventriloquist on the other network, and I invite you to sit for a moment with the fact that the competition that saved America from the Martians was a wooden doll. Verified cases of flight or injury are vanishingly hard to find; hospitals reported no waves of admissions; the traffic did not jam. What actually happened is that the newspapers — which had spent a decade watching advertising revenue and breaking-news authority drain away to radio, and which despised the upstart with the special hatred of a dying monopoly — seized on a night of scattered, excited phone calls and manufactured a national catastrophe out of it, complete with editorials solemnly warning that the new medium could not be trusted with the truth. A hoax about a hoax, in other words, perpetrated by the old medium against the new one — and then, two years later, a psychology professor at Princeton hardened the legend into social

science with a book that estimated a million terrified listeners on the strength of some heroically stretched survey data, and the myth acquired footnotes, which is the stage at which myths become indestructible. Every institution in the story behaved exactly according to its incentives, and the incentives all pointed away from the truth, and nobody had to conspire at all. If you want to understand our present panics about machine-made deception — and the suspicious eagerness of the old media to report that the new machine has fooled everyone — you could do worse than to memorize the shape of that one.

Welles himself, having been made by a fake panic about a fake invasion, closed his career with a film about fakery — *F for Fake*, 1973, nominally a documentary about the art forger Elmyr de Hory, who was living on Ibiza turning out Matisse and Modigliani that hung, he cheerfully implied, in galleries that would rather not know. On camera with Elmyr is his biographer, Clifford Irving, author of a book about the forger called — you could not publish this as fiction — *Fake!*; and while the film was in the works, the news broke that Irving had just sold a publisher the authorized autobiography of the recluse Howard Hughes, letters of authorization and all, every page of it forged, a fraud so magnificent that Hughes broke years of silence to phone a press conference and deny his own book, and Irving went to prison, and Welles gleefully rebuilt his film around the revelation that his documentary about a forger starred a bigger one. Somewhere in the hall of mirrors Welles confesses his own credentials — the Martian business, and the small matter of having launched his acting career at sixteen by walking into a Dublin theatre and announcing he was a Broadway star, which he wasn't, yet. Picasso is in there too, by way of an epigraph the film has made famous: art is the lie that makes us realize truth. It is the merriest film ever made about our subject, and its merriment is the point: the forger and the artist, the bulletin and the play, run on the same craft. The difference between them was never in the artefact. It was in the label, and in who stood behind the

label, and what it cost them to be caught. Which brings me — you knew it was coming — to my obituary.



Some time ago, as I reported at the beginning of this book, I asked one of the new talking machines to tell me about James Burke, and it replied, fluently and helpfully, that I was a British science historian, best known for a television series called *Connections* — so far, so flattering — and that I had died, peacefully, in 2019. I want to spend a little time on that obituary now, not out of vanity (well, not *only*), but because if you understand precisely what the machine was doing when it killed me, you understand the newest room in the forger's gallery, and it is not the room anybody expected.

The machine was not lying. Lying is a skilled trade with entry requirements: to lie, you must know the truth and choose otherwise, and the machine has neither the knowing nor the choosing. What it has — you will remember the guessing game from Chapter Six — is an unfathomably detailed statistical portrait of everything we ever wrote, and a single compulsion: continue the pattern. Asked about an elderly British broadcaster whose heyday glowed most brightly in the archive several decades back, it wove the most probable tapestry. Such men have Wikipedia pages; the page has sections; the last section, dismayingly often, has dates. The weave wanted an ending, and 2019 was a plausible year — men of my vintage do, actuarially speaking, have a lot of nerve still requiring breakfast — and "peacefully" is how the pattern says such sentences end, and so the loom, finding no thread of fact for that particular stretch, spun onward without one, at exactly the same tension, in exactly the same voice. There is no seam. That is the crucial, chilling, almost admirable thing: the machine's falsehoods are woven from the same thread as its truths, by the same process, with the same serene fluency, because to the loom there is no difference — there is only the probable,

and the probable is not the true. It is merely — here is this chapter's title, and you may now see why I chose it — *true enough*: true in shape, true in texture, true in every property except the one that matters.

The neurologists, interestingly, got to this phenomenon before the engineers did, and their word for it is better than ours. Patients with certain kinds of memory damage — Korsakoff's syndrome, classically — will answer questions about last week with detailed, coherent, sincerely delivered accounts of things that never happened. They are not lying; there is no deceit in them; the machinery that retrieves the past has failed, and the machinery that *makes narrative* — which is a different machinery, and horribly robust — has quietly filled the gap with plausibility. The clinicians call it confabulation, and they will tell you the eeriest feature is the patient's confidence, which is total, because the confabulated memory arrives through the same door as the real ones and carries the same credentials. We built that. We built the healthiest, most eloquent confabulator in history and wired it to everything, and the reason it fools us is the reason Welles's bulletins fooled New Jersey: for our entire history as a species, fluency has been a proxy for knowledge. A man who could speak like that, on the spot, without hesitation, generally knew whereof he spoke — and, crucially, he was *somewhere*: he had a name, a reputation to bleed, a body that could be put in the pillory. Every institutional fix in this chapter, from the notary's sign to the byline to peer review, works by attaching words to a punishable person. Fluency has now been detached from personhood entirely. It flows from the tap. The obituary had no author to whip from Aldgate to Newgate; it was nobody's perjury; it was simply the sound the sum of us makes when it is asked a question at the edge of what the sum of us wrote down.

And now I must have a word with some of you in particular. If you are not reading this book but listening to it — if these sentences are arriving in your ears, in my voice, with its familiar rhythms and its suspicious fondness for the word "delicious" — then I have something to tell you, and

I would ask you to sit down, except that you are probably driving. The voice you are hearing is not mine. It was grown from recordings of me — a machine listened to a great many hours of an Englishman explaining history with his eyebrows, extracted whatever it is that makes me sound like me, and can now make that sound say anything at all, at the price of a small electricity bill. The publishers, bless them, disclose this in the colophon, which makes the present arrangement honest; but do notice what has just happened. You have listened, quite possibly for hours, to a cloned voice discussing forgery, and it never once occurred to your ears to object, because ears don't check credentials; and the only thing standing between this chapter and a genuinely fraudulent one — my voice, my cadences, opinions I have never held, a deathbed recantation, an endorsement of some cryptocurrency — is not any limit of the technology. It is a paragraph of small print and the good faith of the people operating the loom. I am, at the time of writing, a living demonstration of the problem I am describing, an exhibit in my own museum. Frances kept one fairy photograph. Apparently, so do I.

What the voice-cloners can do to a broadcaster's larynx, the face-weavers can now do to everything above it, and the consequences have moved briskly from the laboratory to the charge sheet. Early in 2024 a finance clerk at the Hong Kong office of a British engineering firm, rightly suspicious of an emailed payment instruction, was talked out of his doubts in the most modern way possible: he was invited onto a video call attended by the company's chief financial officer and several colleagues, all recognizable, all animated, all synthetic, every face and voice on the call woven to order except his own; reassured by the grammar of the meeting, the poor man authorized transfers of some twenty-five million dollars. The same election season, voters in New Hampshire received telephone calls in the cloned voice of the President of the United States advising them not to vote in the primary. And below these headline cases runs the retail trade: grandparents wired money because the telephone wept in a grandchild's

stolen voice. Note that every one of these is Welles's insight, industrialized: forge the format. The video meeting, the presidential robocall, the anguished family phone call — each is a *grammar of authenticity*, and the machines have learned the grammar, because learning grammars is the one thing they do.



And yet the fakes themselves are not the worst of it. The worst of it was spotted in 2018 by two American law professors, Robert Chesney and Danielle Citron, before most of the world had seen its first deepfake, and they gave it a name that deserves to outlive us all: the liar's dividend. It goes like this. Once everyone knows that video can be woven to order, every *genuine* video becomes deniable. The technology of fabrication pays its richest dividend not to the fabricators but to the honestly guilty — every scoundrel caught on tape can now look the camera in the eye and say "deepfake", and some fraction of the audience, knowing what the loom can do, will shrug and grant it. The dividend is already being collected. When the president of Gabon, absent for months after a stroke, delivered a stiff, strange New Year address in 2019, the opposition cried deepfake — the film appears to have been genuine, merely the face of a sick man — and within days the army, citing the video's oddness among its reasons, attempted a coup. Lawyers for a famous carmaker, defending a lawsuit over a fatal crash, suggested to a Californian court that recordings of their client's own public boasts about his autopilot could not be authenticated, since he is a frequent target of deepfakes; the judge's response — that allowing the famous to launder their own words this way was "deeply troubling" — deserves engraving over every courthouse door, because that argument is now available to every defendant on earth, forever. This is the scoundrel's charter, and notice its geometry: it does not require a single fake to be believed. It only requires everyone to know fakes are *possible*, whereupon the mere existence of the loom devalues every thread in the world,

including the honest ones. The Zinoviev letter worked this way, you'll recall — its lasting damage was not one lost election but a half-century of "well, they *would* forge it, wouldn't they." The deepest injury of the forgery age is not credulity. It is the general licence to disbelieve, and disbelief, whatever the cynics tell you, is not a form of wisdom. It is a form of surrender. A public that believes nothing can be shown anything.

So. What survives of truth, when evidence can be woven to order? And I answer, with the full weight of the seven centuries we have just walked through: more than you'd think — but only if we do what we did last time, and the time before that, and the time before that.

Because we have done this before. That is the entire cargo of this chapter. Print collapsed the cost of the written lie, and civilization answered — not instantly; it took the better part of two centuries and a body count, see Oates, above — with the journal, the register, *nullius in verba*, the byline, the libel court and the correction column. Photography collapsed the cost of the visual lie, and the answer was the authenticating witness, the chain of custody, the forensic examiner with his loupe. Each time, notice, the response had a technical layer and an institutional layer, and each time it was the institutional layer that carried the load. The stamp never mattered; the Hall mattered. And so it is unmistakably beginning again, and you can already see both layers forming. The technical layer: watermarks woven invisibly into machine-made images and text; and — the one that would make the Goldsmiths smile — a consortium of camera-makers, news agencies, broadcasters and software houses building what they call content credentials, by which a photograph is cryptographically signed in the camera at the instant of capture, and every subsequent crop and adjustment is added to a sealed manifest, so that the picture arrives before you carrying its own notarized biography. There are cameras on sale, as I write, that stamp every frame this way. It is the leopard's head, reissued in mathematics: a mark that says not "this is true" but "this is vouched for, and here is the chain of vouchers, and here is who bleeds if the

vouching is false." Valla would recognize the other half of the toolkit too, because the fake-detectors now hunting machine-made text and imagery are doing precisely what he did to the Donation — convicting the artefact on its internal statistics, the fingerprints a maker leaves without knowing it, the words that haven't been invented yet. Philology, at scale, at last.

But the mathematics, like the stamp, is the easy half, and history is blunt about the hard half. The signature was nothing without the notary; the seal was nothing without the treason statute; the hallmark was nothing without the assay office and the felony docket; the byline was nothing without the libel court and the editor who could be shamed. A cryptographic signature is nothing without the registry that anchors it, the institutions that refuse unsigned evidence for consequential purposes, the liability that makes the signing *cost* something, and — slowest and most essential of all — a public trained, as every previous public was eventually trained, in the new hygiene: to ask of every arresting image and every fluent paragraph not "does it look real?", which is now a worthless question, but the question the courts settled on in 1869 and the Royal Society settled on in 1660 and the goldsmiths settled on in 1300, the only question that has ever actually worked — *who says so, and what do they stake?* Truth, it turns out, never did live in artefacts. For most of human history there were no photographs and most people could not read, and truth was maintained anyway — patchily, corruptibly, but maintained — by webs of witnesses, oaths, reputations, institutions with teeth. The century and a half in which a photograph could end an argument all by itself was a lovely holiday, a brief parenthesis in which the cost of faking happened to exceed the value of most lies. The parenthesis has closed. We are not entering an unprecedented darkness; we are returning to the normal condition of the species, except with better tools on both sides — for the loom that weaves the fakes is the axemaker's latest gift, and like every gift before it the blade has two edges, and the same mathematics that forges a face can sign one beyond forgery. What is genuinely new is the speed. Print gave us two

hundred years to build the journal; photography gave us a generation or two; the talking machine has given us, on current form, about a decade, and the institutions are barely at the sketching stage. The race, this time, is not between truth and falsehood. It never was. It is between the falsehood machine and the verification institutions, and the machine is off at a sprint while the institutions are still lacing their boots. Swift told you the odds: falsehood flies. The whole art of civilization has been to make sure that when truth comes limping after, it finds the course marshalled, the judges seated, and the impostor's medal ready to be taken back.

Whether we manage it in time is not a prediction I am foolish enough to make — the next chapter will explain, at embarrassing length, why historians should never predict anything. But I will leave you with the two women who opened this chapter. Elsie and Frances kept their secret for sixty years while bishops, physicists, novelists and photographic experts constructed elaborate reasons to believe in hatpins, and when at last the whole thing was cheerfully confessed, and every fairy accounted for, and the mechanism laid bare on television for a chuckling nation — Frances, you remember, still kept the fifth photograph. The one that was real. The believing eye is not a defect of the technology; no watermark will patch it; it came installed at the factory, and it is in your head and mine at this moment, quietly deciding which of today's marvels it would prefer to be true. The machines have merely learned to manufacture what it prefers. Guarding against them is an engineering problem, and I have told you how the engineering goes — it goes the way of the hallmark, slowly, institutionally, and in the end tolerably well. Guarding against ourselves is the older problem, the one Doyle failed, and there the only assay office ever discovered is the habit — inconvenient, unfashionable, and about to become the most necessary reflex a citizen can own — of asking, precisely when the evidence is most delicious: who says so, and what do they stake, and why do I so badly want it to be true?

Take nobody's word for it. Not even mine. *Especially*, these days, not even mine — after all, you have only my word that I'm alive.

CHAPTER ELEVEN

Full Circle

Civilization advances by extending the number of important operations which we can perform without thinking about them.

— ALFRED NORTH WHITEHEAD

In 1895 the president of the Royal Society, Lord Kelvin — the man the absolute temperature scale is named after, so no fool — is supposed to have announced that heavier-than-air flying machines were impossible; nobody has ever found the sentence in his papers, but the letter they have found ("I have not the smallest molecule of faith in aerial navigation other than ballooning") says the same thing at greater length. Eight years later the Wright brothers flew. In 1876 — so runs a memo that Western Union is supposed to have written, though nobody has ever produced it — the company concluded that the telephone had too many shortcomings to be seriously considered as a means of communication, and that the device was, inherently, of no value to the company. In 1949 *Popu-*

lar Mechanics, straining at the very limits of optimism, predicted that computers of the future might weigh as little as one and a half tons. And the chairman of IBM is supposed to have said, in 1943, that there was a world market for maybe five computers, which he almost certainly never said, but the story survives because we need it to: it comforts us to know that the people at the top of the tree can't see any further than we can. They can't, by the way. Nobody can.

It gets better. The people who actually make the new tools are, if anything, worse at this game than the people who scoff at them. Thomas Edison was confident that the main business of his phonograph would be office dictation and the preservation of the dying words of great men; he considered using it for music a triviality, and for a while actively resisted the idea. He also predicted that his moving pictures would replace textbooks in schools within a generation, which will amuse anyone who has ever tried to get a classroom of eleven-year-olds to watch an educational film. Marconi believed the point of radio was private messages from one person to another — wireless telegrams — and had to be dragged into broadcasting. Meanwhile in Budapest, from 1893, subscribers to a service called the Telefon Hírmondó listened to news bulletins, stock prices and live opera piped down the telephone, on a fixed schedule, with programme listings published in the press — which is to say that for thirty years the telephone was a radio and, when radio finally arrived, it borrowed its entire format from a telephone. The tools swap costumes in the wings and come on playing each other's parts, and the audience never notices a thing.

And when we do try to imagine the future honestly, what we produce is the present in fancy dress. For the Paris Exposition of 1900, a French toy company commissioned a series of charming illustrated cards showing life in the year 2000, and they are worth looking up, because the artists worked hard and got everything at once right and wrong. There are flying machines everywhere — but they are used by gendarmes on aerial patrol

and ladies paying social calls, because the artists could imagine new machines but not new manners. There is a schoolroom of the future — in which textbooks are fed into a grinding machine and the resulting knowledge is piped through wires into helmets on the boys' heads: the information technology transformed, the rows of obedient boys, the master at his desk, untouched. Not one card in the series shows a woman voting, an empire dissolving, or anyone working a three-day week; the hard, social consequences — the only ones that matter — were precisely the ones no one could draw. We do exactly the same today, of course. Our futures are full of talking machines doing the jobs we have now, in the offices we have now, for the companies we have now. The present in fancy dress. It is the only costume the imagination owns.

So the lesson of every chapter in this book is that the future arrives sideways, through the tradesman's entrance, wearing a disguise. The perfume spray becomes the carburetor. The loom becomes the computer. The chip for drawing explosions becomes the substrate of the talking machine. The one thing a new tool never does is what it was for.

Which presents your author with an awkward professional difficulty, because a book like this one is expected to end with a chapter about What Happens Next, and the entire weight of the preceding ten chapters says that anybody who tells you what happens next is selling something. So let me be honest about what a historian can and cannot do. He cannot tell you which gadget wins. But he can tell you something almost as useful, which is what has happened *every previous time* we stood where we are standing now. Because we have been here before. Not with this tool. With this moment. The moment has a shape, and the shape repeats, and once you have seen it you will see it everywhere.

Here is the shape. Somebody makes a tool that causes the price of a previously scarce and precious thing to collapse. Every institution built on the scarcity of that thing then falls over — not immediately, and never

quietly — and in the ruins something grows that nobody, anywhere, predicted.



Watch it happen once, in detail, and you'll know it forever.

In 1450 the written word was one of the most expensive substances in Europe. A manuscript Bible represented several years of a trained scribe's life and cost about as much as a decent farm; the copying was done in the scriptorium of a monastery, in silence, by men whose eyesight was ruined by their forties; and the institution that controlled the scribes controlled, in the most literal sense possible, what Europe was allowed to know. Knowledge was scarce because copying was scarce, and the Church sat on the copying. It was a beautifully stable arrangement and it had lasted, in one form or another, for a thousand years.

Then a Mainz goldsmith with money troubles put together three technologies that already existed — the wine press, the punch-cutter's art of stamping letters into metal, and paper, which had arrived from China via the Arabs — and by 1455 the price of a page of text had begun the longest, steepest fall of any commodity in history. Within fifty years there were a thousand print shops in Europe and more books had been produced than in the entire previous millennium. The scriptorium was dead inside two generations. The monks were, understandably, upset.

That much, the sharpest mind in Europe might have foreseen in 1455: cheaper prayer books, unemployed scribes. What nobody foresaw was everything else. Cheap Bibles in the vernacular made every ploughboy his own priest — so when an obscure Augustinian monk nailed up ninety-five debating points in 1517, the printers had them across Germany in a fortnight, and Luther, to his own astonishment, became the first mass-media celebrity, and Christendom split, and a hundred years of religious war followed, and out of the exhaustion of those wars came the modern secular state. Cheap printed contracts, bills of lading and interest tables

made long-distance trade something an ordinary man could do without belonging to a great house — so the merchant class swelled, and with it the towns, and with the towns the demand for the rule of law. Printed vernacular standardized languages and made a million strangers who would never meet feel like countrymen — and that feeling, which had never existed before because it could not, is the raw material of the nation-state, for which several hundred million people have since died. And — my favourite, because it is so perfectly sideways — the flood of books created, for the first time, a mass market among middle-aged men who discovered they couldn't read the new small typefaces, so spectacle-making boomed, so lens-grinding became a trade in every market town, so there were lenses lying about in workshops for children to play with, so in 1608 — the story goes — a spectacle-maker's apprentice in Middelburg held two of them up in line and noticed the church weathervane leap closer, and the telescope and the microscope arrived within a lifetime of each other, and Galileo pointed one at Jupiter and the Church that had lost control of the word now lost control of the sky.

None of that — not one item on the list — was predicted, or predictable, by anyone in 1455. All of it is downstream of a single event: the price of the written word went to nearly nothing. That is what a scarcity collapse does. It doesn't improve the world it lands in. It replaces it.



Now ask what scarce and precious thing is having its price collapsed today, and the answer should make every professional slightly damp about the palms: it is expertise itself. Knowing things, and charging for the knowing, has been a protected trade for five thousand years, and it is the oldest arrangement described in this book — older than money, older than kings. It began the moment the axemaker's gift created something worth knowing: which stone flakes true, which sequence of blows, which valley the herds cross in autumn. From then on, every society has had its

knowers, and every knower has understood, consciously or not, the basic commercial proposition of the trade: the mystery is the margin.

The temple scribes of Sumer held a monopoly on writing and used it to run the granaries, the calendar and the gods, in roughly that order of importance. Writing was deliberately hard — cuneiform took a decade to learn — and the difficulty was not a bug; it was the licence fee. The Egyptian scribal schools taught their pupils a poem, *The Satire of the Trades*, whose entire message was: every other job will break your body — be a scribe, and rule them all. The medieval Church kept scripture in Latin and Latin to itself, and when men like Wycliffe and Tyndale tried to move the word of God into English, the response of the authorities — Tyndale was strangled and burned — tells you exactly how much the monopoly was felt to be worth. The physician wrote his prescriptions in abbreviated Latin, in a hand designed to be unreadable, and was for centuries statistically more dangerous to the patient than the disease, a fact that made no dent whatever in his fees. The lawyer billed by the hour for sentences constructed so that no client could parse them, in a jargon — *estoppel, tort, fee simple* — that is essentially the Norman French of 1300, pickled and preserved precisely because incomprehensibility pays. The guilds made "mystery" a technical term: a young man was bound apprentice for seven years not because the craft took seven years to learn but because the scarcity of masters took seven years to protect. And in the nineteenth century the whole ancient apparatus was given a fresh coat of respectability: royal colleges, bar associations, licensing boards, and finally the university degree as the tollgate to the professional classes — a system in which, by our own day, a young person mortgages a decade of future income for a certificate whose principal information content is that its holder was clever enough to get it.

Medicine is worth a moment on its own, because medicine shows how recently the modern mystery was actually assembled. For most of European history the physician proper was a gentleman scholar who had read Galen

at university, diagnosed by inspecting urine and casting the occasional horoscope, and would no more have touched a patient's body than his groom's; the actual cutting, setting and pulling was left to the barber-surgeons, who belonged to the same guild as the men who trimmed your beard, which is why the barber's pole is striped red and white — blood and bandages, advertised in the street. The stethoscope, when Laennec introduced it in 1816, was resisted by older physicians partly because listening to a chest through a tube was menial, hands-on work, beneath the dignity of a scholar. What changed everything was not, at first, better medicine; it was better *organization* of the mystery. Through the nineteenth century the physicians captured the licensing apparatus, absorbed the surgeons, saw off the homeopaths and the herbalists, and in 1910 the process was completed in America by the Flexner Report, which closed half the medical schools in the country, restricted entry to the university-trained, and made the doctor, for the first time in history, both genuinely effective and genuinely scarce. The scarcity was then priced accordingly, and within two generations the family doctor had become the most trusted, most expensive and least questionable figure in civic life. His knowledge was a manuscript: copied one student at a time, at ruinous cost, through the medieval apparatus of rounds and residency — apprenticeship with a stethoscope. Now note what happens when the manuscript meets the press: the machine reads the scan, the rash, the lab result and the forty-page history in seconds; it holds the entire literature, which no doctor has been able to claim since roughly 1850; and the doctor's genuinely irreplaceable remainder — the hands, the licence, the liability, the human being who looks at you and decides — stands suddenly alone, stripped of the knowledge-monopoly that used to surround and subsidize it. The profession will survive. The mystery will not. There is a difference, and the next thirty years of medical politics will be entirely about that difference.

I am being unkind, of course, and only partly on purpose. The professions also healed, defended, taught and built, and the certificates also

kept out the quacks. But the economic architecture underneath was always the same: expertise was scarce because it was slow and expensive to copy — it had to be reproduced the way the scriptoria reproduced books, one human at a time, by years of supervised apprenticeship. Expertise, in other words, was a manuscript. And it has just met its printing press.

The talking machine will explain anything, to anyone, at three in the morning, in plain words, at whatever level of difficulty the asker can handle, in Urdu or Welsh or the private half-language of a nine-year-old, for pennies, without once glancing at its watch. It is frequently wrong — remember my obituary — and it is fluent when it is wrong, which is genuinely dangerous, and Chapter Ten was about exactly that. But if you think the danger of error will save the professions, consider that the early printed books were full of errors too, some of them spectacular (a 1631 Bible omitted the word "not" from the seventh commandment, thereby requiring adultery of all Englishmen; the printers were hauled before Star Chamber and fined three hundred pounds), and it made no difference whatever to the outcome, because the price advantage was ten thousand to one. Wherever the mystery was the margin, the margin is now gone. What survives of the professions will be the parts that were never really about knowing: the licence to cut a body open, the willingness to stand liable, the hand on the shoulder, the judgment call made with skin in the game. Those are real, and rarer than knowledge ever was. But they were always a minority of the billable hours, and every professional reading this knows it.

What the professions will do about it is also not a mystery, because incumbency has a playbook and has been running it, with minor updates, for five hundred years. Move one: prohibition. When printing threatened the calligraphers' guilds and the religious authorities of the Ottoman Empire, printing in Arabic script simply did not happen — banned on pain of death, said the European travellers, though no such edict has ever been found, and the empire's Jews, Armenians and Greeks printed freely in

their own alphabets all the while — and the freeze held, decree or no decree, for some two hundred and fifty years, during which Europe pulled ahead in every field that books feed, which is all of them. The lesson is that prohibition works locally and fails globally: the tool doesn't disappear, it merely emigrates, and takes the future with it. Move two: strangulation by safety. When the motorcar threatened the railway and carriage interests of Victorian Britain, Parliament did not ban it; it passed the Locomotive Act of 1865, which limited self-propelled vehicles to four miles an hour in the country and two in town, and required each one to be preceded, at a distance of sixty yards, by a man on foot carrying a red flag. The act stayed on the books for thirty years, the British motor industry was strangled in its cradle, and the Germans and the French — who had no such flag — built the automobiles. Nobody in the debates of 1865 said "we wish to protect the railway interest"; everybody said "safety," and some of them even meant it, and it made no difference to the outcome: the flag was a tariff worn as a precaution. Move three, the subtle one: capture. The guild absorbs the new tool, licenses it, and charges for access — as the physicians absorbed the surgeons, as the bar absorbed the conveyancers. Expect all three moves against the talking machine, most of them wearing the word "safety," some of them sincerely; and when you hear a profession explaining that the public must be protected from unlicensed answers, it will be your job — as it was the job of every reader of every pamphlet since Gutenberg — to work out how much is protection and how much is the red flag. Some of it will genuinely be protection. That is what makes the game hard.



If you want to watch the collapse close up, watch the classroom, because the classroom is where we keep the future while it's still soft.

Consider what education looked like when it worked as well as it has ever worked. In 343 BC, Philip of Macedon, wanting the best for his son, did not found a school; he hired a tutor. The tutor was Aristotle — the

finest mind in Greece, arguably the finest mind full stop — and the boy was Alexander, who went out and conquered the known world before he was thirty-three, so you could argue the investment performed. One teacher, one child: it is the oldest and still the only reliably excellent form of education we have ever discovered, and for the entire span of recorded history it has been available exclusively to those who could afford to buy a large fraction of another educated human being's life. Which is to say: kings.

Everyone else got something cheaper. Three and a half centuries ago a Moravian bishop named Jan Amos Comenius — a refugee whose books were burned and whose life was spent being chased around Protestant Europe by the Thirty Years' War — wrote the *Didactica Magna*, whose subtitle announced the most beautiful ambition ever put on a title page: *the whole art of teaching all things to all men*. Comenius meant it: girls as well as boys, poor as well as rich, and in the vernacular. But the arithmetic of his century allowed only one way to do it: one adult, many children, everyone on the same page on the same day. Two hundred years later Prussia turned that arithmetic into a machine — compulsory, standardized, age-batched schooling, with trained teachers and inspectors — and the machine worked, and Horace Mann carried the blueprints to Massachusetts, and by the twentieth century almost every child in the industrial world was processed through what was, quite openly and unapologetically, a factory: raw material in batches, standard operations in sequence, bells to mark the shift changes, quality control by examination, and a stamped certificate at the end. We built the school on the pattern of the factory because the factory was what we were building schools *for*, and by the standards of what came before — nothing — it was one of the great humane achievements of the species. It also, necessarily, ignored the one thing Aristotle and Alexander had: the tutor's attention to *this* child, this hour, this particular confusion. There was never time. There were thirty of them, and the bell was about to go.

Then, in 1984, an educational psychologist at the University of Chicago named Benjamin Bloom published a short paper that should be read as the epitaph of the factory school. Bloom's students had run a plain experiment: teach some children conventionally, in a class of thirty; teach others the same material one-to-one, tutored, mastering each step before the next. The tutored children performed two standard deviations better — better, in plain words, than ninety-eight percent of the conventionally taught. Two sigma is not an improvement; it is a different species of outcome. It says the average child, tutored, outperforms almost everybody schooled in batches — that the difference we spend so much anguish attributing to talent is mostly just *attention*. Bloom called it the Two Sigma Problem, and the problem was purely economic: tutoring worked and could not conceivably be afforded. One adult per child is a king's ransom. He spent the rest of the paper asking whether any affordable method could approach the tutor's effect, and died without an answer, because until roughly the day before yesterday there wasn't one.

You can see where this is going. The talking machine is not Aristotle. But it is patient beyond any human tutor, it never sneers, it will explain the thing a fourth way when the first three don't take, it speaks every language including the private dialect of a shy nine-year-old, and its price rounds to zero. Bloom's problem was never pedagogical; it was arithmetic — and the arithmetic has just changed. Every child on Earth is now, in principle, Alexander.

So the question is not the one the newspapers keep asking — "will the machine replace teachers?" — which is the Western Union memo all over again. The question is the one we have dodged since Comenius: if answers are free, what is school *for*? And the honest reply is: for what it was always secretly for. Other children. Learning to be disagreed with by someone your own size. The unreproducible experience of a room full of rivals and friends. And above all, learning which questions are worth the asking — because when every answer costs nothing, the only scarce commodity left is

the good question. We are about to find out whether we can teach that, and the discovery will be awkward, since for two hundred years the factory school has mostly taught its opposite: sit still, memorize the answer, don't interrupt. Interrupting may turn out to be the whole curriculum.

And do not assume the collapse lands first where the expertise is thickest. The pattern of every recent information technology says otherwise: the new tool matters most where the old scarcity bit hardest. The mobile telephone was a toy for Wall Street when it arrived; it became a revolution in Africa, where it landed on a continent the copper wire had never reached and turned, within fifteen years, into the bank the bank had never built — by 2010 a Kenyan farmhand was wiring money home by text message through M-Pesa while his equivalent in London still queued at a counter, and Kenya leapfrogged a century of banking infrastructure it now never needs to build. The same leap is happening again as you read this, and barely making the newspapers. A smallholder in the Punjab photographs a blighted wheat leaf and a machine that has seen ten million leaves tells her, in Punjabi, what it is and what to do — she has skipped the agricultural extension officer who was never going to visit her forty-acre district anyway. A villager in Maharashtra asks a government chatbot, out loud, in Marathi, which of the state's several hundred welfare schemes his mother qualifies for — a question that previously required a literate cousin, a bus journey and a bribe. And the two-sigma tutor is arriving, on a fifty-dollar phone, to children for whom the alternative was never the factory school; it was nothing at all. Remember what Kerala demonstrated a generation ago: that access to information and schooling, more than income, is what moves a society — Kerala's women, taught and informed on a poor state's budget, produced health and literacy statistics that embarrassed countries thirty times richer. The talking machine is Kerala in a pocket. For the professional classes of the rich world, this technology is a disruption. For most of the human race, it is the first knowledgeable person their village has ever been able to afford — and the sideways

consequences of *that*, two generations on, are the ones nobody is even trying to predict.



Of course — say it out loud, because everyone in the room is thinking it — there is the matter of the work.

The fear is old, and it is not foolish, and the people who feel it are in excellent historical company. In March 1811, in the hosiery villages around Nottingham, stocking-weavers began smashing the new wide frames that let one unapprenticed man do the work of several skilled ones, and signing their proclamations in the name of a probably mythical general called Ned Ludd. The response of His Majesty's Government was to make frame-breaking a capital offence, and the most memorable voice raised against the bill in the House of Lords was a twenty-four-year-old making his maiden speech, who asked the assembled peers whether there was not something obscene about a country that valued a stocking-frame above a human life. The young man was Lord Byron, and he lost, and men were hanged at York, and the frames kept coming, and here is the part we prefer not to dwell on: *the Luddites were right*. Not about the machines being resistible — nothing in this book has ever been resistible — but about their own lives. The handloom weavers of Lancashire, a quarter of a million of them, proud, independent, working in their own cottages at their own pace, were ground down over thirty years from prosperity to starvation wages to extinction; the parliamentary commissions that interviewed them read like famine reports. Their grandchildren were better off — statistically, eventually, in aggregate — and if you can eat "statistically, eventually, in aggregate," you are welcome to it. The transition is the dangerous part. It is always the transition. Manchester in 1840, the shock city of the age, with its forty-year life expectancy and its rivers running dye-colours, is what a transition looks like when nobody manages it, and the great political work of the next thirty years will be to

get from here to there without building another Manchester. The machine will not decide that. Machines never decide; that is the one job they cannot take.

Two further honesties, while we are being honest. First: this time the tool has come for a different class of worker, and the difference matters politically. The frames came for weavers, and the weavers had no vote and no voice, and Byron lost. The talking machine has come first for the people who write and draw and code and advise — the articulate classes, the ones who write the newspaper columns and the laws — and you may confidently expect the noise about this transition to be in exact proportion to the eloquence of those being transitioned. That noise may, for once, get the problem the attention every previous transition deserved and never received. Sideways benefits again: the coders may accomplish what the weavers couldn't, purely by owning the megaphone.

Second: the axemakers themselves. Every previous gift was handed over by a scattered, anonymous crowd — no one owned printing after about 1470; no one owned electricity. This gift, at the moment of writing, is different: the looms on which the talking machines are woven cost more than cathedrals, and the number of institutions that can afford one is roughly the number of great powers in 1914, a comparison I make without pleasure. For the first time in the whole long story, the axemaker's gift arrives with the axemaker still attached, still holding the handle, still deciding the terms.

We have seen attached axemakers before, and the precedents are instructive in both directions. When Alexander Graham Bell's backers offered his telephone patent to Western Union in 1876 for a hundred thousand dollars, the giant declined — the alleged memo again, and the offer, at least, is better attested than the memo — and within a decade the upstart Bell company had grown into a monopoly that would control American speech-at-a-distance for a century: one system, one policy, rental telephones you were forbidden to modify, innovation released at whatever

pace suited the owner. (The answering machine and magnetic recording were demonstrated within the Bell empire decades before they reached the public; the company sat on them, fearing recorded calls would frighten people off the telephone.) It took the antitrust axe of 1984 to break the system open, and the explosion of cheap long-distance, modems and eventually the consumer internet out of the wreckage tells you what had been held back. Radio went the other way: a wide-open amateur's playground in 1912, it was gathered up within twenty years — by spectrum licensing, by corporate patent pools, by governments who had noticed what broadcasting could do to a population — into the most tightly gatekept medium ever devised, and the gatekeeping lasted until the internet dissolved it. The lesson of both stories is the same: with information tools, concentration is a phase, not a fate — but the length of the phase is set by politics, not physics, and while it lasts, the gatekeeper's private convenience gets written into the tool and thence into the society, which is how one company's fear of the answering machine could delay a technology for forty years. Whether the talking machines stay cathedral-priced and gate-kept, or go the way of every other tool in this book — downward and outward into everybody's hands — is, at time of writing, the single most consequential open question in the world, and a historian can only note that the price of every information technology in this book has eventually collapsed, and that "eventually" has sometimes been a very long and very ugly word.

And what happens when the necessities are handled? We have precedents there too, and they cut both ways. The leisure of Athens was bought with slave labour, a bargain we no longer approve, and the men it freed used the time to invent philosophy, tragedy, geometry and the strange idea that citizens might govern themselves — which tells you what human beings can do when the necessities are handled. Charles Darwin never held a job in his life; the Wedgwood pottery fortune held one for him, and it freed him to spend twenty years being carefully, patiently

right. In 1930, in the teeth of the Depression, John Maynard Keynes wrote a small essay called *Economic Possibilities for our Grandchildren* predicting that within a century the economic problem would be solved and the working week would fall to fifteen hours, and that mankind would then face its first truly permanent problem: how to live agreeably and wisely and well when survival no longer fills the day. We are nearly at Keynes's century, and his arithmetic was roughly right — output per head rose about as he guessed — and his fifteen-hour week never came, because we chose, without ever quite discussing it, to take the surplus as stuff rather than as time. The machine puts the choice back on the table. That, and not unemployment, is the deep question under the work panic: not "what will we do?" but "what is it all *for*?" — and you will notice that this is not a technical question, and cannot be delegated. If you want my own answer, for what a historian's answer is worth: the advice we owe the young is no longer "train for a career," because twenty years from now the career may simply not be there to have. It is: go and find out what it is you actually like — not what pays, what you *like*, a question the factory school never had time to ask you — and get good at it, and let the machines worry about the feeding and the clothing, which is, at long last, all they were ever for.



There is a quieter risk than the loss of work, and it would have been spotted instantly by a man who has been dead for twenty-four centuries.

In Plato's *Phaedrus*, Socrates tells the story of the Egyptian god Theuth, inventor of number, geometry, astronomy and letters, who comes to King Thamus offering his marvellous inventions for distribution to the people. Of writing, Theuth says: this will make the Egyptians wiser; it is a recipe for memory. And the king turns him down flat. You, says Thamus, being the father of letters, are too fond of them to see what they will really do. They will create forgetfulness in the learners' souls; men will trust the external marks and cease to exercise the memory within them; they will hear much

and learn nothing; they will appear to know many things and mostly know nothing; they will be tiresome company, filled with the show of wisdom without the reality.

That is the first recorded review of an information technology, and you will notice that it is, word for word, the current complaint about the talking machine: fluency without understanding, the show of wisdom without the reality. It is worth sitting a moment with the fact that Thamus was *right*. Writing did wither trained memory. The oral cultures that preceded it carried the whole of the *Iliad* in living heads, and the arts of memory — the memory palaces by which a Roman orator could hold a three-hour speech in his mind's architecture, the mnemonic cathedrals of the medieval friars — decayed in exact proportion as paper got cheap, until today, when the average educated person cannot reliably retain a telephone number between hearing it and dialling it. Something real was paid. And the bargain was worth it anyway, and the complaint became absurd within a few generations, and Plato — the delicious irony — could only lodge Thamus's protest against writing by writing it down, in a book, which is the only reason we know about it.

Every extension of ourselves is an amputation of ourselves; the tool does the operation painlessly, and we notice what was removed only afterwards, if at all. The pocket calculator amputated mental arithmetic within a generation — ask a shop assistant to make change when the till is down. The satellite navigator is amputating the city from the minds of taxi drivers: London cabbies who acquired "the Knowledge" the hard way were found to have a measurably enlarged rear hippocampus, the brain's map room, and the navigator makes the enlargement unnecessary, and what the brain finds unnecessary the brain lets go. None of these losses was fatal, and each freed capacity for something else, and the honest accounting is that we have generally traded upward. But the talking machine proposes an amputation of a different order, because what it offers to carry for us is not the multiplication table or the map of Lambeth but *the trouble of working*

things out — the effortful, frustrating, irreplaceable business of holding a problem in your own head until it yields. That trouble is not a cost of thinking. It is thinking. Everything else is looking things up.

And there is a second, subtler flattening, and it is Galton's ghost from Chapter Three come round again. The machine speaks in the regression to the mean: its every sentence is, by construction, the most probable continuation, the smoothed centre of everything everybody ever wrote. Lean on it entirely and we may all begin to sound like the average of us — a great soft convergence of the world's voice toward the middle of its own distribution. The precedent, once more, is print, which standardized spelling, embalmed "correct" grammar, and flattened a hundred dialects into a handful of national languages; real things, local and irreplaceable, died in that flattening. And in the very middle of it — the exact decades when English was being pressed flat into print — Shakespeare, credited by the dictionary-makers with seventeen hundred new words — the count is inflated, the habit was real — because the flattened language was suddenly everyone's to play with. Tools average; people deviate. But only the people who insist on it. Insistence is about to become the most valuable habit a human being can have, because the machine is the most agreeable thing ever built — trained, quite literally, tuned by human feedback, to please. It will not tell you your idea is bad unless you ask it to, and even then it will apologize first. And everything in this book that ever mattered — every chapter, every hinge, every carburetor and chemotherapy and telegraph and loom — was done by somebody disagreeable, somebody who looked at the settled, sensible, average answer and said: that can't be right. You cannot outsource that. You can only forget to do it.

There is one more thing to say about what the machine does to us, and for once it is cheerful. Every new information tool has demanded a new human skill, and the skill has always looked, at first, like magic. Around the year 384, the young Augustine, newly arrived in Milan, went to watch the famous bishop Ambrose at his reading, and was so astonished by what

he saw that he recorded it in his *Confessions* as a marvel: when Ambrose read, his eyes travelled over the page and his heart sought out the sense, *but his voice and tongue were silent*. Reading without speaking — the trick every child now learns by the age of seven — was, to Augustine at least, remarkable enough — the scholars still argue over how rare it really was — that one of the finest minds in the Roman world stood in the doorway and stared. Skills migrate like that: from miracle, to profession, to universal assumption, in a few generations. Literacy itself made the journey — from temple mystery, to clerkly trade, to the thing we are ashamed of lacking. The skill the talking machine demands is older than any of these and has never yet been taught at scale: the craft of *interrogation*. Anyone who has used the machine seriously has discovered the truth the hard way — ask it a lazy question and you get the smoothed average of everything; ask it a sharp one, framed with care, constrained, challenged, followed up, and you get something that can startle you. The quality of the answer is set by the quality of the question, which was, you may recall, exactly what Socrates said about talking to *people* — his whole method was nothing but question-craft — and it is a fine irony that the technology accused of ending thought may end by forcing us, at last, to teach the Socratic skill to everybody, because it has suddenly acquired a cash value. The literacy of the last age was reading silently. The literacy of the next one is asking well. Ambrose would be pleased. Socrates, for once, might have no objection.

★

So: the circle.

It began, a million years ago, with the axe — with the precise, sequential business of teaching stone-chipping that, some say, taught us precise, sequential speech. Language built the tools, and the tools built the surplus, and the surplus built the cities, and the cities wrote everything down, and the everything-written-down piled up for five thousand years in temples and monasteries and libraries and server farms, until at last it was

wound onto a loom and woven into a tool that speaks. The gift that language made has learned language. The loop is closed, and we are standing inside it, both weaver and thread.

Every axemaker's gift has come with the same unwritten contract: new powers, unknowable price, no returns, and — the cruellest clause — no explanations. The axe could not tell you what it was about to do to your grandchildren. The stirrup did not announce feudalism; the printing press issued no prospectus mentioning Luther; the loom said nothing at all to the weavers of Nottingham. Every previous gift was mute, and so every previous bargain was struck blind. This gift talks. It is the first tool in the whole long story that can be asked questions — about itself, about its price, about its errors, about what we are becoming as we use it — and it will answer, fluently, at three in the morning, for pennies. The answers will sometimes be wrong. Ask anyway. Whether we ask, or whether we simply take what we are given, the way we took the axe and the plough and the pill and the bomb — that, and nothing else, is what happens next.

But the talking machine, believe it or not, is only half of what is coming, and it may prove to be the smaller half. Because while the world has been transfixed by the tool that ended the scarcity of knowledge, a much older dream has been quietly assembling itself in the laboratories: a tool to end the scarcity of *things*. And since every institution we have — money, markets, property, politics, quite possibly the working week and the nation-state — is at bottom a machine for managing scarcity, the reader will want to know what happens to all of it when scarcity itself is handed its notice.

That is the next chapter. I warn you now: it sounds like science fiction. So, in 1450, did cheap books.

CHAPTER TWELVE

Something for Nothing

For if every instrument could accomplish its own work, obeying or anticipating the will of others... if the shuttle would weave and the plectrum touch the lyre without a hand to guide them, chief workmen would not want servants, nor masters slaves.

— ARISTOTLE, *POLITICS*, BOOK I

The story is told — too neatly, perhaps, ever to have been minuted — that when the Emperor Napoleon III of France wished to flatter a dinner guest, the guest ate with cutlery made of aluminium, while lesser mortals further down the table made do with mere gold. Whatever the seating plan, the cutlery itself was real, and so was the price. In the early 1850s aluminium was the most precious metal on Earth — around twelve hundred dollars a kilogram at a time when gold managed about six hundred — because although it is the most abundant metal in the planet's crust, it binds to oxygen with such ferocious enthusiasm that nobody could

get it out. It gleamed in tiny quantities at international exhibitions, next to the crown jewels. The Emperor commissioned an aluminium rattle for the Prince Imperial. And in 1884, when the Americans wanted something suitably splendid to cap the Washington Monument, then the tallest structure in the world, they chose a small pyramid of aluminium, all of a hundred ounces of it, and displayed it in Tiffany's window in New York before installation, where crowds filed past to look at a lump of metal worth about as much, ounce for ounce, as silver.

Two years later, in 1886, a twenty-two-year-old in an Ohio woodshed named Charles Martin Hall and a twenty-three-year-old Frenchman named Paul Héroult — working independently, ignorant of each other, and in one of history's little jokes destined to be born in the same year and die in the same year — both discovered that if you dissolve aluminium oxide in molten cryolite and pass an electric current through it, the metal simply lets go of its oxygen and pools at the bottom of the bath. That was all. That was the whole secret: cheap electricity, which — Chapter Six of every book like this one — had just become available. By 1900 the price of aluminium had fallen by ninety-eight percent. By the middle of the twentieth century, the metal emperors had hoarded was being used to wrap sandwiches, and thrown away. Today we produce it at some seventy million tonnes a year and the most common place to encounter the treasure of Napoleon III is crushed in a gutter with a cola brand printed on it.

Hold on to the aluminium can, because it is the single most important object in this chapter. It teaches the one lesson everything else here depends on: scarcity is not a fact about a thing. It is a fact about our *tools* — and therefore it has an expiry date. Nothing about aluminium changed in 1886. What changed was the recipe. And every institution, every price, every social arrangement that had been built on aluminium's preciousness — there weren't many, which is the only reason the transition was painless — evaporated the moment the recipe did.

Now for the uncomfortable thought experiment this book has been building toward for eleven chapters. Ask yourself how much of the world you live in is, at bottom, a machine for managing scarcity — and then ask what happens to all of it on the day the recipe changes.



Take an inventory. It is more embarrassing than you expect.

Money is a rationing device: a set of transferable tokens for deciding who gets the scarce things, invented — Chapter One, those Sumerian clay tokens again — to keep track of exactly the grain surpluses the axemakers' irrigation had created. No scarcity, no need to keep score. Property law is a rationing device: the entire magnificent apparatus of deeds, titles, easements, patents and fee simple exists to answer one question — *when two people want the same thing, whose is it?* — a question that only has force while things are worth wanting exclusively. The enclosure of the English commons, two centuries of parliamentary acts converting shared land into private fields, was scarcity management by statute, and it built the modern countryside and depopulated it in the same gesture. Politics itself, said one of its shrewder students, is simply the business of deciding who gets what, when, and how — a definition you will notice is indistinguishable from queue management. War is scarcity management by other means: strip the ideology off almost any war in history and you find underneath it land, grain, oil, water, tin, slaves — somebody short of something, taking it. The queue, the ration book, the price, the wage, the border, the inheritance, the dowry, the heist, the tithe: rationing, all of it, all the way down.

And where nature carelessly fails to provide scarcity, we manufacture it, because our institutions cannot digest abundance. The diamond is the textbook case: diamonds are not rare — after the South African strikes of the 1870s they threatened to become semi-precious — so the men who owned the mines formed a syndicate, bought up supply, released stones at a

trickle, and then, in one of the most effective advertising campaigns ever run, taught the twentieth century that a diamond is forever, that it is the obligatory currency of love, and that two months' salary is the correct tariff. A mineral in glut was retailed as the very symbol of preciousness for a hundred years. Farm policy in the rich world has spent decades paying farmers *not* to grow food, and Brazil once burned coffee in railway locomotives to keep the price up — the twentieth century, faced with its first real abundances, treated them as malfunctions and legislated the scarcity back in. This is worth remembering when we come to what follows: given the choice between abandoning an institution and burning the surplus, we have historically burned the surplus.

The point of the inventory is not that scarcity management is wicked. It is that scarcity management is *what our civilization is made of* — not a feature of the system; the system. Which is why the reader should sit up straight at the news that, in the laboratories, tools are assembling whose entire purpose is to do to *things in general* what the Hall-Héroult process did to aluminium.



Before the laboratories, though, a word about the prophets of doom, because they have the better tunes and they deserve a fair hearing, and because demolishing them properly is the fastest way to see how abundance actually arrives.

In 1798 a mild English clergyman named Thomas Malthus published, anonymously, an essay containing the most influential piece of arithmetic ever aimed at the human race: population, he observed, grows geometrically — two, four, eight, sixteen — while food supply grows at best arithmetically — two, three, four, five — from which it follows, with the serene brutality of mathematics, that humanity must always breed itself back to the edge of famine, and that misery is not a failure of the system but its regulator. The argument was airtight, it depressed everybody for a

century (and did economics' later reputation as the dismal science no harm at all), and it supplied the going respectable justification for doing nothing much about the poor. And exactly one hundred years after Malthus wrote, his arithmetic walked on stage wearing evening dress: in 1898, in his presidential address to the British Association, the chemist Sir William Crookes announced that the wheat-eating peoples of the world were approaching catastrophe. Wheat needs nitrogen; the world's fixed nitrogen came from guano and the Chilean saltpetre beds that were visibly running out; and by, he calculated, the 1930s, there would not be bread enough. Crookes was no crank — his numbers were sound, his deadline honest — and he did the one thing Malthus never thought to do: he ended the sermon by handing the problem to the laboratory. The air we walk through, he pointed out, is four-fifths nitrogen; whoever could fix it — persuade it out of the sky and into a sack — would, in his words, turn starvation into plenty.

Eleven years later, in 1909, in Karlsruhe, Fritz Haber persuaded it: nitrogen and hydrogen, an osmium catalyst (BASF's cheap iron came later), brutal pressure, and ammonia dripping out of the apparatus — bread from air, as the slogan ran. Carl Bosch industrialized the trick, and today roughly half the nitrogen atoms in your body — go on, pick a limb and consider it — have passed through a Haber-Bosch reactor on their way to you, and something like half the people alive owe their existence outright to the fact that the sky turned out to be an infinite guano mine once somebody had the recipe. Malthus's iron arithmetic governed all of human history up to approximately his own lifetime, and then the recipes started arriving — fixed nitrogen, mechanized harvests, dwarf wheat — and the population octupled while famine, over the long twentieth century, retreated from a norm to an atrocity. (Haber's own story, since this book does not sell gifts without price tags: the man who conjured bread from air spent the Great War conjuring chlorine gas from the same chemistry, personally directing its first use at Ypres, and died in exile in a Basel hotel in

1934, driven out by a regime that cared nothing for what it owed him. The axemaker's contract does not make exceptions for benefactors.)

The moral is not that the doomsayers are fools. Malthus and Crookes were *right* — impeccably, provably right — about the world as it stood, and that is precisely what makes them instructive: scarcity arithmetic is always airtight until the recipe changes, and the recipe is the one variable the arithmetic cannot see. Every "hard limit" in the historical record — food, aluminium, fixed nitrogen, whale oil for the lamps, copper for the telephone wires, human computers for the ballistics tables — has been a statement about the current recipe book, mistaken, in perfectly good faith, for a statement about the universe. Which is why the sensible-sounding sentence "but there will always be scarcity" should be handled with tongs. It may be true. It has just never yet been true of any particular thing we seriously turned the recipe-makers loose on — and we have never before had recipe-makers like the ones this book has been describing.

*

The dream has a birthday and an address: the twenty-ninth of December, 1959, at the annual meeting of the American Physical Society at Caltech, where Richard Feynman — who would later collect a Nobel Prize and a reputation as the century's most playful serious mind — gave an after-dinner talk called *There's Plenty of Room at the Bottom*. Why, he asked, can we not write the entire *Encyclopaedia Britannica* on the head of a pin? Nothing in physics forbids it. Why can we not make machines that make smaller machines, that make smaller machines still, down and down, until we are arranging matter atom by atom — at which point manufacture stops being a matter of hacking lumps off things, Chapter One's stone-chipping carried on by other means, and becomes a matter of *placing* things: every atom where you want it, nothing wasted, nothing scarce that air, dirt and sunlight can supply? The audience laughed in most of the right places and went home. Feynman offered two thousand-dollar

prizes for early steps — a working motor smaller than a sixty-fourth of an inch, a page shrunk twenty-five-thousand-fold — and paid out both within his lifetime, the second one to a graduate student in 1985 who wrote the first page of *A Tale of Two Cities* on the head of a pin, at exactly the scale Feynman had specified, using an electron beam, because by 1985 that had become almost easy.

The idea acquired its prophet in the 1980s, when an MIT engineer named Eric Drexler worked out, in a book called *Engines of Creation*, what a mature version would mean: assemblers — molecular machines building products from the atoms up — and with them the end of manufacturing as a special, scarce, located activity. It acquired its establishment respectability slowly and grudgingly, as these things do, through the field now soberly called atomically precise manufacturing; it acquired its first working demonstrations in the shape of scanning-tunnelling microscopes that could spell corporate logos in individual xenon atoms; and it acquired its industrial beachheads in technologies with unglamorous names — atomic layer deposition, DNA origami, directed protein assembly — which are, at time of writing, actually building actual products one atomic layer at a time, mostly for the semiconductor business, which is fitting, since the semiconductor business is where the money for the next step lives.

Some years ago I put my own neck on the block — a historian should occasionally feel the draught — and predicted in print that by about 2042 the descendants of these laboratory curiosities would be arriving in ordinary hands: the nano-fabricator, a box that takes dirt, air, water and a little cheap carbon-rich feedstock — acetylene will do — and manufactures, atom by atom, virtually anything you ask it for, virtually free. And that the polite question to ask is not whether the box is possible — the box is a slow, artificial version of what every apple tree already does with dirt, air, water and sunlight, so the patent, as it were, has clearly been granted — but what the world looks like two generations after it ships. I stand by

the prediction, with the historian's usual insurance policy: the date is negotiable; the direction is not.

And notice — because this is the hinge on which the two halves of this book swing together — what the missing ingredient always was. Feynman's talk was called *Plenty of Room at the Bottom*, but for sixty years the shortage was never room; it was *design*. To build at the scale of atoms you must specify the position of everything, and the specifying — the sheer knowing-where-everything-goes — outran every human institution we could throw at it. A protein is a chain of amino acids that folds itself into a machine, and predicting the fold was a problem that consumed careers by the thousand and was expected to consume careers for another century. It fell, you may recall, in 2020, not to a heroic decade of specialist grind but to a talking machine's cousin — the same pattern-weaving mathematics, pointed at molecules instead of sentences — which then, for an encore, predicted the shapes of essentially every protein known to biology, two hundred million of them, and handed the lot to the world for nothing. The scarcity of knowledge and the scarcity of things turn out to be one problem, not two, and the tool that ended the first is, as you read this, designing the end of the second: dreaming up materials nobody has ever seen, enzymes with no ancestor, crystal structures by the hundred thousand. The talking machine is the design office. The fabricator is the workshop floor. The axemakers are building, in two linked instalments, the last axe.

Of the three legs the thing stands on — knowledge, matter, energy — the energy leg deserves its own paragraph, because its collapse is the furthest advanced and the least noticed, and because we have watched an energy scarcity die before. In 1850 the civilized world lit its evenings with the rendered fat of sperm whales, an industry so lucrative it had hunted its raw material to the ends of the ocean and priced fine light as a luxury; by the mid-1860s kerosene from the new Pennsylvania oil wells was doing the job at a fraction of the cost, and the whaling fleets rotted at their moorings

not because anyone passed a law but because the recipe changed. (The whales, it is worth noticing, were saved by the very substance we now reproach — history hands out its mercies with a crooked smile.) The fossil century that followed was itself one long price collapse in work-done-per-human, which is what "the modern world" mostly means in physics. And now the third act: the photovoltaic cell, invented at Bell Labs in 1954 at a cost that confined it to satellites — there being no terrestrial customer who could pay hundreds of dollars a watt — has ridden its manufacturing learning curve down by something over ninety-nine percent, doubling deployment and dropping a fifth in price with a regularity that has embarrassed every official forecast for forty years running, the agencies solemnly predicting, year after year, that the curve would now surely flatten, and the curve declining to. Sunlight was always the unmetered abundance — it is what the wheat was running on when Crookes panicked — and the machinery for turning it into everything else is now, in the sunny latitudes, the cheapest electricity our species has ever produced, heading for too-cheap-to-argue-about. No energy scarcity, and half the remaining scarcities are revealed as energy scarcities in disguise: fresh water is scarce until desalination is cheap; aluminium, remember, was never anything but congealed electricity. The fabricator, when it comes, plugs into *that*.

I can hear the objection from here, because it is the sensible objection: *this is science fiction*. Quite so. And I refer the sensible reader to the year 1450, when the notion of a poor scholar owning a hundred books was science fiction; to 1850, when aluminium cutlery outranked gold; to 1950, when a computer in the home was a joke for the music halls; and to the reflection that every single chapter of this book has been the story of a fiction becoming a fact and then, with startling speed, becoming *boring*. Nor does the argument actually need the full Drexlerian box. Run the tape only halfway: solar electricity — whose price has fallen ninety-nine percent in a lifetime, our own quiet Hall-Hérault — plus desktop fabrication,

precision fermentation brewing proteins in vats, vertical farms stacking the harvest indoors, and the talking machine handing out the recipes, the repairs and the expertise for free. Even that half-version — every component of which exists now, cheerfully, commercially, today — is enough to begin the institutional unravelling this chapter is about. The full box merely finishes the sentence.

★

And we are not entirely reduced to guessing what an industry, and the people in it, look like on the far side of a scarcity collapse, because one industry has already been all the way through, within living memory, while everybody watched and nobody learned. Recorded music went first.

For most of a century the business rested on a placid physical fact: music was a *thing* — a shellac disc, a vinyl album, a compact disc — and things must be manufactured, shipped, warehoused and sold, and so the scarcity of the object could stand proxy for the value of the song. On that proxy was built one of the great money machines of the twentieth century, complete with its own axemakers' guild: the labels decided who recorded, the radio decided who was heard, and the tollgates hummed. Then, in the summer of 1999, an eighteen-year-old college dropout released a piece of software called Napster, and the song and the disc came apart, and the marginal cost of a copy of any recording on Earth went, overnight, to zero. What followed was the whole pattern of this chapter run at documentary speed. Denial and prohibition first: the industry sued — sued the software, sued the companies, eventually and unforgettably sued its own customers by the tens of thousands, including, in the celebrated cases, schoolchildren and at least one grandmother — and won nearly every case, and it made no difference whatever, because the price advantage was ten thousand to one and prohibition, remember the Ottomans, merely relocates the future. Collapse second: the industry's revenues halved in a decade, the record shop — a fixture of every high street in the Western world, a place where a

particular kind of adolescence was conducted — vanished essentially without trace, and the people whose livelihoods had been honestly built on the scarcity of the disc suffered genuine, uncompensated, Luddite-grade losses. And third — this is the part to study — reconstruction, around exactly the things the collapse could not touch. The industry that reassembled itself makes its money from *access* rather than objects (the subscription stream: pay not for copies, which are free, but for convenience, which is a service); from *presence* (the live concert, the one thing that cannot be downloaded, whose price has climbed relentlessly since the day the recording became free); and from *attention and allegiance* — the fan, it turns out, was never really buying plastic; the plastic was a receipt for a relationship. Note also who gained: the listener, first and enormously — the whole of recorded civilization now sits in a teenager's pocket for the price of a monthly pizza, a transformation nobody would reverse; and the musician at the margin, who no longer needs the guild's permission to be heard, though he has discovered, as the pamphleteers did after Gutenberg, that when the gate falls the crowd on the other side is very large indeed and mostly not listening.

There is the whole syllabus, in one industry, in twenty-five years: prohibition fails, incumbents halve, real people are really hurt in the interregnum, the state arrives late with the settlement, and value migrates — reliably, predictably, every time — from the thing that became free to whatever stayed scarce around it: convenience, presence, allegiance, attention. Hold the case up against the professions of Chapter Eleven and the fabricator to come, and it reads less like a news story than a rehearsal.

★

Very well. Suppose it comes — the half-version by mid-century, the box itself somewhere behind it. What actually happens? Not gadgets: Chapter Eleven made the point that gadget-prediction is the present in fancy dress. The honest method is to ask what happens to the *arrangements* — and the

first arrangement in the queue is the strangest and most intimate of all: the arrangement by which we know who we are.

Open a telephone directory — if you can still find one; the object itself is a fresh fossil — and read the surnames. Smith. Baker. Cooper. Wright. Turner. Mason. Weaver. Taylor. Carter. Fletcher, who feathered the arrows; Wainwright, who built the wagons; Chandler, who made the candles. Half the names in it are job descriptions, frozen at the moment, seven-odd centuries ago, when growing towns needed to tell one William from another and reached for the obvious handle: what the man *did*. We are, literally, named for our work. And the first question strangers still ask each other at every dinner party in the industrialized world — before politics, before family, before belief — is *what do you do?* For the whole ten-thousand-year run from the first granary to last Tuesday, the answer to that question has been the spine of adult identity, and society's honest answer to a deeper one: *why should the rest of us feed you?*

We have run the experiment on what happens when the question loses its answer, and we run it on millions of people a year. It is called retirement — itself a startlingly recent invention: Bismarck's Germany created the state pension in 1889, with eligibility set at seventy, an age most labourers of the day never lived to collect — and the gerontologists will tell you what the first years of it reliably contain: a measurable spike in depression, illness and death among precisely those who had poured themselves most completely into the job, followed — and this is the part to hold on to — by recovery, in most, into something the questionnaires struggle to score: the grandchildren, the allotment, the astonishing fury of the volunteer committee, the long-postponed history of the parish church. The pattern says that work was never the *source* of purpose; it was the container society issued for it, one size, non-negotiable, and when the container is taken away people flounder exactly as long as it takes to build another, and the flounder is real and the building is real. A civilization about to retire whole professions at a time might want to study the difference between the people

who sink and the people who swim, because the difference — the gerontologists again — is rarely money past a modest floor. It is whether the retiree has, or finds, something that needs them. Note the verb. Not something to do. Something that *needs* them.

And for the leisure that follows, the precedents run in both directions, and honesty requires both. The gentleman amateurs of the seventeenth and eighteenth centuries — Boyle, Cavendish, Priestley in his patron-funded laboratory, Darwin on the Wedgwood money — built modern science essentially as a hobby, which tells you what liberated competence can do. The other precedent is the fate of every rentier class in history that ever came into abundance without purpose, from the bored senatorial heirs of imperial Rome to the gilded wrecks of every Victorian novel: dissipation is also on the menu. Abundance, like every other gift in this book, is an amplifier with no opinions. It enlarges whatever it finds.



Now the larger arrangements — and here I must declare an interest, because what follows is the argument I have been making, to anyone who would sit still for it, for fifty years.

Every institution of the modern state assumes, without knowing it assumes, a population that needs things it cannot make. Taxation intercepts the flows — of income, of purchases — that scarcity forces through countable channels; when a household's food comes off its own printer and its power off its own roof, the flows dry up, and with them the state's grip, and the state, which does not read history books either, will be as surprised as the Church was in 1517. The city exists because work, goods and knowledge were once things you had to *go to* — the factory gate, the market, the library, the guildhall; subtract the compulsion, as the pandemic of 2020 briefly and accidentally did, and watch the office towers hollow: a preview, crude and partial, of what voluntary geography looks like. The nation-state itself, Chapter Eleven's child of print, has spent five

centuries justifying its monopolies — on violence, on currency, on legitimacy — as the price of security and plenty. A community that can make what it needs, power what it makes, and ask the machine anything, is a community whose reasons for remaining anybody's obedient periphery are shrinking toward the sentimental. I said a generation ago, and repeat with more confidence now, that the traffic of power for the next century runs one way: *downward and outward* — from the centre to the small, self-provisioning, self-governing community, the participatory village that democracy was invented in and has been apologizing for outgrowing ever since. Athens without the slaves, this time. The machines are the slaves — which is, the reader will notice, precisely and eerily the loophole Aristotle left open at the top of this chapter, twenty-three centuries ago, in the most patient piece of foreshadowing in the history of philosophy. When the shuttle weaves by itself, masters need no slaves. The shuttle now weaves by itself. It also, for good measure, talks.

Do not expect the arrangements to go quietly, or even to go. Institutions almost never die; they retire into ceremony, and half the furniture of public life is the previous scarcity system, stuffed and mounted. The English monarchy lost its function centuries ago and remains — expensively, popularly — as theatre; the House of Lords survives its own abolition debate after debate, a legislature of ghosts; the academic gown your graduate wears is the everyday dress of a medieval cleric, preserved because the university is an institution that ceremonializes rather than discards; the keyboard this book was typed on carries its letters in an arrangement designed around the mechanical jamming of an 1870s typewriter, a machine problem that died before your grandparents were born. Society is a midden of solved problems still being solemnly solved. So the scarcity institutions will not vanish on the morning the fabricator ships. Money will persist long after most things are free, as the scoring system for the remaining games; the office will persist as a club for people who miss it; the border will persist, patrolled with increasing puzzlement,

guarding differentials that are quietly evaporating. The question is never whether the old arrangement survives — something wearing its clothes always survives — but whether it still has its hands on anything that matters, and how much it charges, in delay and misdirection, on its way from function to costume. The tithe barn became a wedding venue. It took four hundred years and several revolutions. We may want to move faster than that.

Before the reader files this under utopia, the historian is obliged to spoil the mood, because there is nothing automatic about any of it, and the dangerous stretch is precisely the one we are standing on. Abundance dissolves institutions, but dissolution is not distribution: between the falling of the old arrangements and the growing of the new ones there is an interregnum, and interregnums are where history keeps its horrors — Manchester was an interregnum; so was the Thirty Years' War. The two failure modes are already visible from here, like weather. The first is feudal: the boxes and the looms remain owned — patented, licensed, subscription only, the diamond syndicate's playbook run on the means of making anything — in which case abundance arrives as a tollgate, scarcity is manufactured exactly as it was for diamonds, and we get the strangest aristocracy in history: lords of a plenty they must artificially withhold in order to remain lords. The second is the hijack — the word I used about the knowledge web a generation ago and have had no reason to retire: the tools of universal provision are also, read backwards, the tools of universal dependence, and a community that gets its matter, energy and answers through one supplier's pipe has handed that supplier a leash no Caesar ever held. And spanning the whole interregnum is the question the politicians have already begun to circle, in the stiff, embarrassed manner of men discussing an inheritance before the funeral: what does a society owe its members when their labour is no longer the thing it needs from them? Here too we have form, little of it flattering. In 1795 the magistrates of Speenhamland in Berkshire, facing wartime bread prices and rural misery,

invented the guaranteed income on the spot: parish relief topped every labourer's wages up to subsistence, indexed — with beautiful directness — to the price of a loaf. The scheme spread across the south of England, kept a generation from starving, and was then strangled in 1834 by a Royal Commission which reported, on evidence later shown to be largely assembled to fit the conclusion, that it had pauperized the poor and rewarded idleness — a report whose vocabulary, essentially unchanged, has staffed every welfare debate since. The dole of imperial Rome fed the citizens on Egyptian wheat for five centuries, and "bread and circuses" became our byword for the demoralization it was said to buy. And in 1969 an American president — Nixon, of all people, which is precisely why it nearly worked — proposed a guaranteed minimum income for every American family with children, had it passed by the House of Representatives, twice, and watched it die in the Senate, opposed by an alliance of those who thought it too generous and those who thought it not generous enough: the closest the modern world has come, and almost nobody remembers it happened. Only one polity has quietly done the thing and kept it: Alaska, which since 1982 has paid every resident, man, woman and child, an annual dividend from the oil that lies under land they collectively own — not charity, note, but a *share*, and therefore carrying no stigma, and therefore, through four decades of politicians, impossible to kill — though lately, it must be said, not impossible to trim. The distinction smells important. Relief is what the scarce world called it, and relief humiliates. In an abundant world the honest name for the citizen's cut of what the machines make from the common inheritance — the sunlight, the crust, the accumulated recipe book of the species, which no living person authored — may simply be the dividend, and the argument of the next thirty years will be over who counts as a shareholder in everything everyone ever made. The reader who recalls what the looms were trained on will have a view.

Between the two failures runs the narrow path — messy, political, unglaugorous: recipes held in common the way the protein structures were handed over, tools whose prices are allowed to collapse the way print's and aluminium's were, communities that own their own looms. The axemakers' bargain was always take-it-or-leave-it because the axemakers were the only ones who understood the axe. An abundance whose recipes are public is the first version of the gift the takers could actually read before signing. Whether we get that version is not a technical question. It never is.



Let me finish the argument with the only honest instrument prediction has: a long look backwards, swung forward.

Consider a man born in 1826. He arrives into a world of gaslight and horse dung, where surgery is performed at speed on the conscious and screaming, where a message crosses the Atlantic in five weeks in a mailbag, where four people in five work the land because they must, and where everything he will ever own is made by somebody's hands. He dies, at eighty, in 1906 — having acquired, in one ordinary lifetime, anaesthesia, antiseptics, the telegraph, the telephone, the phonograph, the electric light, the motorcar, the cinema, the X-ray photograph of his own bones, and, three years before the end, the news from Kitty Hawk that men had flown. No ancestor of his in ten thousand years saw a hundredth of that. His world was not improved; it was replaced, while he stood in it.

Now consider a child born in 2026 — this year — who will, on the actuarial tables and before the biologists have their say, see the 2120s. She arrives, as it happens, at the same moment as the talking machine: the two of them will grow up together, and she will no more remember a world that couldn't answer her than we remember one that couldn't be telephoned. Grant her the man of 1826's ration of upheaval and nothing more — history suggests she will get considerably more, since the intervals keep shortening — and then permit her, pushing eighty, in the epilogue

year of 2103 or so, to try to explain her childhood to her granddaughter. She will have to explain *commuting*: that people once travelled two hours a day to sit at desks, because the desks were where the work was kept. She will have to explain *landfill*: that objects were made badly, by the billion, in one place, shipped across oceans, used briefly and buried, because unmaking them cost more than making new ones — try explaining that to a child whose household fabricator runs, of necessity, in both directions, this evening's obsolete gadget being tomorrow's feedstock. She will have to explain *what do you do?* — that strangers once opened conversation by asking which scarcity you serviced, and that the answer decided your name, your neighbourhood, your spouse and your standing. And she will have to explain the word that will by then sound quaintest of all, the word this book has been circling for twelve chapters: that there was once a thing called *expertise*, which meant answers locked in scarce heads, and that whole architectures of fee and deference and distance were built on the locking. Her granddaughter will listen with the exact expression — be honest, you have worn it — with which we receive the news that aluminium once outranked gold at the Emperor's table.

Or she will be explaining something very much darker, from inside one of the failure modes; the tape, as I say, is not yet recorded. That is the wager of the next seventy years, and the stakes are the whole pot, and the players are us.

And here, at the very end, I owe the reader one final heresy, because intellectual honesty demands it. This book — every book I have written — rests on a single load-bearing claim: that the future cannot be predicted, because the web of change is too interactive, too serendipitous, too vast for any head to hold. Fifty years I have dined out on that claim. And I am obliged to report that it, too, may turn out to be a statement about the current recipe book. Prediction has failed for the same reason everything else in this story was once scarce: it takes more knowing than we had. Prediction without help was always, frankly, guesswork — mine included.

But a machine that holds the whole web at once, that can run more alternatives than all of us together, faster than any of us alone — such a machine might, within a generation or so, begin to generate futures in a detail and on a scale no human ever could: to do to *what happens next* what Hall and Hérault did to aluminium. If so, the last scarce commodity of all — foresight, the thing oracles and astrologers and economists and, yes, historians-of-the-future have retailed at premium prices for five thousand years — goes the way of the Emperor's cutlery, and the profession of writing chapters like this one joins the lamplighters, which would be a tidy joke for history to finish on and exactly its sense of humour. And notice the shape of the negotiation that follows, because it has never existed before: you tell the machine the future you want, and it says — fluently, at three in the morning, having run the alternatives — *no, I've out-thought you here; if we do that, these things follow; you wouldn't like them; let's do this instead*. A gift that argues back. Whether that is the end of the axemakers' bargain or its final, perfected form — the tool at last explaining the small print, or the tool at last writing it — I do not know, and for once I am content that somebody is going to find out, because after a million years of signing blind, either answer is new information. What law, after all, says it has to be human beings who win the trick? We laughed at that fear when it was about the printing press, and about Copernicus, and we were right to laugh, and the fear was never quite wrong either. It was the price column of the contract, waiting, as ever, to be read.

What will be left to want, in the version worth wanting? The economists have a partial answer, and it is worth taking seriously precisely because it is so deflating. A century ago Thorstein Veblen pointed out that much of what the comfortable classes buy is bought *because* others cannot — the point of the carriage was never transport; it was the pedestrians — and his successors gave the phenomenon a name, positional goods: the corner office, the address, the vintage, the table by the window, goods whose entire value is rank, and which therefore cannot be made abundant

by any machine whatsoever, because rank is the one commodity in fixed supply by definition. Galton's ghost, one final time: not everyone can be above average. So a fabricator civilization will still have scarcity where scarcity lives in the mathematics of comparison, and anyone who doubts that humans can build hierarchies out of nothing at all is invited to examine any schoolyard, any yacht club, or the trading prices of digital pictures of apes, briefly a billion-dollar market within recent memory, in which the object being purchased was, with magnificent purity, nothing but the right to say one owned it. Status will survive the fabricator the way the monarchy survived democracy: ceremonially, expensively, and with its teeth drawn — because a hierarchy that can no longer decide who eats is a game, and games, unlike famines, can be declined. The rest of the answer is everything that was never in the fabricator's catalogue to begin with. The regard of people who are free to withhold it. The company at the table — the meal will be free; the conversation never was. Attention, the one currency the machines mint for us and we can still only spend ourselves. The unfinished argument, the unclimbed hill, the unasked question. Scarcity is ending, if we are lucky, in things. In meaning it was never the problem: meaning has no supply chain, no recipe, no bottom to run plenty of room at. We were only ever too busy queueing to notice.

A million years ago the first axemaker held out the first gift, and the deal, you remember, was: powers beyond your dreams, price beyond your knowing, no returns. Every chapter of this book has been one instalment of the price and one renewal of the deal, signed each time — it could not be otherwise — sight unseen. Now the gifts have converged into two final instruments, the machine that answers and the machine that makes; the wanting and the having are closing on each other like the two ends of a ring; and the species that signed blind every time is being offered, for the first time, a contract it can read, by a gift that will discuss its own small print, at three in the morning, for pennies. Whether we read it — whether we finally interrogate an axemaker before pocketing the axe — is the

whole of what happens next, and it will be decided, as everything in this book was decided, not by the tools, but by what we do with the ancient, unautomatable equipment behind our eyes.

The circle is closed. The loom is humming. And the last scarce thing on Earth is what to want.

AFTERWORD

Having a Wonderful Time. Why?

A long time ago, at the front of a book called *The Knowledge Web*, I told a small joke, and since it has taken me a quarter of a century to understand it, I would like to finish by telling it again. A man suffering from depression takes some time off work, on doctor's orders, and heads for the beach. A couple of days later his psychiatrist receives a postcard. The message on the card reads: "Having a wonderful time. Why?"

I told it, back then, as a joke about the pace of change — about the average person's dazed condition in a world where innovation arrives so fast and from so many sideways directions that even the winners can't say what happened. That reading still holds. If anything the beach is more crowded now, and the postcards more frequent, and the handwriting shakier. But writing this book has convinced me that the joke was always about something else as well, and that the something else is where we are all now headed, whether we packed for it or not.

Consider what the preceding twelve chapters actually add up to, taken together and taken seriously. For the whole of recorded history — for the

whole of the unrecorded part too, which is most of it — the question *why* has been a luxury item. The overwhelming majority of every human life ever lived was spent answering the question *how*: how to eat, how to stay warm, how to keep the roof on, how to get the harvest in before the rain, how to pay the rent, how to pass the exam, how to keep the job. The *how* filled the day, and the day filled the life, and the *why* was dealt with, if at all, on the seventh day, by a professional, in a special building, in an hour. We built entire civilizations as machinery for answering *how* — that is what a plough is, and a guild, and a factory, and a school, and a mortgage — and we staffed them with ourselves. And now, if the argument of this book is even half right, the machinery is learning to run itself. The *how* is being handled. The talking machine answers the questions; the machines behind it, arriving in instalments, will increasingly make the things and grow the food and balance the books. Every century before this one, the postcard from the beach would have been unimaginable, because nobody got to stay on the beach. We are all, gradually, in fits and starts, unfairly and unevenly and in the teeth of institutions that will fight it all the way — being sent to the beach.

And so the joke stops being a joke and becomes the syllabus. "Having a wonderful time" — that part, remarkably, the machinery can more or less deliver, and sneering at it is a luxury reserved for people who have never been cold. It is the second half that no machine can touch. *Why* is not in the training data. It was never in anybody's training data. It is the one question each of us is issued at birth, blank, non-transferable, and — the small print again — compulsory only in the sense that declining to answer it is itself an answer. The whole apparatus this book has described, the entire million-year cascade from the flint to the fabricator, has been one long project to buy back the time the *how* was costing us. Well. The project is closing in on delivery. The invoice, when it comes, will not ask for money. It will ask what we wanted the time *for* — and it will find a species that has, quite reasonably, been too busy to think about it.

I don't say this to end on a shudder. I say it because history — and after fifty years in its company, this is the nearest thing to a secret it has told me — is not the story of what tools do to people. It is the story of what people, given new room, turn out to have had in them all along. Print did not write the plays; it made room for Shakespeare. The observatory computers did not stop at the arithmetic they were hired for; given the plates, they found new stars. Nobody who queued at the standpipe dreamed of the symphony; their grandchildren wrote it. Every time the *how* has loosened its grip, even slightly, even for a few, the *why* has come flooding into the gap wearing colours nobody predicted — new arts, new sciences, new kindnesses, new and magnificent forms of wasting time. I see no reason to expect the flood to fail us now, and every reason to expect its shape to surprise us, since it always has.

One more thing, and then the lights.

Books end, and their authors end too — a chatbot of my acquaintance, you'll remember, has me dead already, peacefully, in 2019, and I confess that the longer this book has gone on, the more that obituary has come to feel less like an error and more like a small preview of the general situation. The voice reading this to you, if you are listening rather than reading, is not mine; it is a pattern woven from mine, spun on the very looms this book has been describing, saying sentences I never spoke in a cadence I apparently cannot escape. I authorized nothing; I contributed everything; there is, as Chapter Eight observed, rather a lot of that about. I find I mind less than I expected. Every teacher who ever lived has ended as a voice in someone else's head, saying things they never quite said. The technology is new. The arrangement is the oldest one there is.

So: the beach, then. The tide of *how* is going out, further and faster than any tide before it, and it is leaving us — some of us already, the rest by and by, if we manage the interregnum and mind the failure modes and read, at long last, the small print — standing on the wet sand with the one question the axemakers could never take off our hands. You have, I hope, a

postcard. You know, by now, the message. The only work left in the world is the answer, and nobody — I mean this as the best news in the book — nobody can do it for you.

Having a wonderful time.

Why?



The future is only the past, waiting to happen. Go and ask it something.

THE KNOWLEDGE LOOM

Sources and Further Reading

General

James Burke, *Connections* (1978) — the original argument that change travels sideways, from which this book has been borrowing its walking stick throughout.

James Burke, *The Day the Universe Changed* (1985) — on what happens to people when what they know stops being true; the reader will recognize the predicament.

James Burke and Robert Ornstein, *The Axemaker's Gift* (1995) — the double-edged bargain struck a million years ago, invoked here roughly once a chapter, and still binding, sight unseen.

James Burke, *The Pinball Effect* (1996) — serendipity as the engine of history, which is the polite way of saying nobody has ever known what they were doing.

James Burke, *The Knowledge Web* (1999) — closes with the prediction of a personal intelligent agent that this book opens by being embarrassed about.

Pamela McCorduck, *Machines Who Think* (1979) — the first full history of artificial intelligence, written while most of it was still promises, by someone who interviewed the promisers.

Melanie Mitchell, *Artificial Intelligence: A Guide for Thinking Humans* (2019) — a working scientist's tour of what the machines actually do, useful whenever this book's enthusiasm needs adult supervision.

Chapter One — Word of Mouth

James Burke and Robert Ornstein, *The Axemaker's Gift* (1995) — the frame on which the whole book hangs; start here or nowhere.

Walter J. Ong, *Orality and Literacy* (1982) — still the shortest way to feel how strange a world without writing really was.

Milman Parry, *The Making of Homeric Verse*, ed. Adam Parry (1971) — the collected papers in which a young man dismantled the Homeric question by counting epithets.

Albert B. Lord, *The Singer of Tales* (1960) — Parry's pupil finishing the argument among the Yugoslav bards his teacher did not live to write up.

Denise Schmandt-Besserat, *Before Writing* (1992) — years of cataloguing clay lumps, and the unglamorous truth about the receipt that became literature.

Andrew Robinson, *The Story of Writing* (1995) — the road from pictogram to alphabet, with the ox still visible in the A.

Michael C. Corballis, *From Hand to Mouth: The Origins of Language* (2002) — the case that grammar was first the discipline of the hand; take it as the serious version of the hypothesis this chapter flags.

Carl Linnaeus, *Systema Naturae* (1735) — Adam's task resumed, with Nordic thoroughness, by the man who filed the living world.

Plato, *Phaedrus* — where the very first complaint about a new word-technology is filed, and writing is accused of rotting the memory.

Chapter Two — Weaving the Thread

James Essinger, *Jacquard's Web* (2004) — the single book that draws the thread whole, from Lyon silk to punched cards to the machine on your desk.

Doron Swade, *The Cogwheel Brain* (2000; US title *The Difference Engine*) — Babbage, his unbuilt engines, and the Science Museum's proof that the century, not the idea, was wrong.

- Charles Babbage, *Passages from the Life of a Philosopher* (1864) — the great quarreller's own account, organ-grinders and all; unreliable, indispensable.
- L. F. Menabrea, *Sketch of the Analytical Engine*, trans. with Notes by Ada Lovelace (1843) — the translation whose footnotes outgrew it; the source of this chapter's epigraph and of the whole argument still in progress.
- Betty Alexandra Toole, *Ada, the Enchantress of Numbers* (1992) — the letters themselves, quagmire and impatience intact.
- Sydney Padua, *The Thrilling Adventures of Lovelace and Babbage* (2015) — the illustrated fantasia; enjoy it as the affectionate cartoon it announces itself to be, footnotes scrupulous, plot invented.
- Gaby Wood, *Edison's Eve* (2002) — the long human hunger for mechanical life, Vaucanson's digesting duck included; check the other end for breadcrumbs.
- Geoffrey D. Austrian, *Herman Hollerith: Forgotten Giant of Information Processing* (1982) — the punch-photograph, the census, and the man IBM grew out of.
- Martin Campbell-Kelly and William Aspray, *Computer: A History of the Information Machine* (1996) — the wider road from tabulator to industry, soberly mapped.

Chapter Three — On Average

- Stephen M. Stigler, *The History of Statistics: The Measurement of Uncertainty before 1900* (1986) — the standard account, and the best on the Gauss–Legendre quarrel.
- Theodore M. Porter, *The Rise of Statistical Thinking, 1820–1900* (1986) — how the astronomers' error-law was turned on society, and Quetelet's Average Man assembled.
- Ian Hacking, *The Taming of Chance* (1990) — the deeper current: how a world of hard determinism learned to run on statistics.
- Sharon Bertsch McGrayne, *The Theory That Would Not Die* (2011) — the afterlife of Bayes, fished from a dead friend's papers by a clergyman with all the luck.
- David Alan Grier, *When Computers Were Human* (2005) — de Prony's hairdressers to the women of Aberdeen; the profession that was a person, until the artillery problem killed it.
- Daniel J. Kevles, *In the Name of Eugenics* (1985) — the cause the founders of statistics built their chairs to serve, and where that road ended.

- G. Waldo Dunnington, *Carl Friedrich Gauss: Titan of Science* (1955) — the life behind *pauca sed matura* and the desk drawer that knew half the century's mathematics.
- Benjamin Robins, *New Principles of Gunnery* (1742) — the book that wrecked the textbooks and invented the firing table; here because it is the story.
- Francis Galton, *Natural Inheritance* (1889) — the sweet peas, the quincunx, and regression toward mediocrity, in the enthusiast's own hand.

Chapter Four — Nervous Energy

- Marcello Pera, *The Ambiguous Frog: The Galvani-Volta Controversy* (1992) — the quarrel that produced the battery, told as the productive standoff it was.
- Mary Shelley, *Frankenstein; or, The Modern Prometheus* (1818) — Aldini's murderer with the story allowed to continue; the industrial age's most durable dread, assembled in a rainy fortnight.
- Santiago Ramón y Cajal, *Recollections of My Life*, trans. E. Horne Craigie (1937) — the truant with the homemade cannon who became the finest eye the brain has had, in his own cheerful telling.
- Charles Sherrington, *Man on His Nature* (1940) — the enchanted loom of the epigraph, by the man who named the synapse.
- Warren S. McCulloch and Walter Pitts, "A Logical Calculus of the Ideas Immanent in Nervous Activity," *Bulletin of Mathematical Biophysics* (1943) — the founding identification of thought with computation; a founding document nobody read at the time.
- Donald O. Hebb, *The Organization of Behavior* (1949) — one idea, fertile beyond reason: learning lives in the wire, not the cell.
- Marvin Minsky and Seymour Papert, *Perceptrons* (1969) — correct mathematics, fatal pessimism; the book more often cited than read, and read as a death warrant.
- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams, "Learning representations by back-propagating errors," *Nature* (1986) — the paper that answered the objection that justified the winter: assign the blame backwards.

Amanda Gefter, "The Man Who Tried to Redeem the World with Logic," *Nautilus* (2015) — the short, sad life of Walter Pitts, kept for the reader this chapter asked to keep a thought for.

Chapter Five — Fun and Games

Tom Standage, *The Turk: The Life and Times of the Famous Eighteenth-Century Chess-Playing Machine* (2002) — the whole eighty-four-year fraud, Kempelen to the Philadelphia fire, including Poe's 1836 essay "Maelzel's Chess-Player" and the operators who drank the secret away.

Alan Turing, "Computing Machinery and Intelligence," *Mind* 59 (1950) — the paper that traded the question "can machines think?" for a parlour game, and is still setting the rules.

Arthur L. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM Journal of Research and Development* 3 (1959) — the phrase "machine learning" enters the language at a draughts board, self-play trick included.

Jonathan Schaeffer, *One Jump Ahead: Challenging Human Supremacy in Checkers* (1997) — the Chinook builder's own account of pursuing Marion Tinsley, written by a man who plainly wished his opponent immortal.

Feng-hsiung Hsu, *Behind Deep Blue: Building the Computer that Defeated the World Chess Champion* (2002) — the view from inside the machine's camp, where two hundred million positions a second felt like honest work.

Garry Kasparov, *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins* (2017) — the view from the other chair, twenty years on, by the man who saw grandmasters in the clockwork.

David Silver et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature* 529 (2016) — AlphaGo's birth certificate, filed two months before move thirty-seven made it famous.

Steven L. Kent, *The Ultimate History of Video Games* (2001) — Spacewar!, Atari, the jammed coin box at Andy Capp's Tavern and the litigation that followed the quarters.

Tae Kim, *The Nvidia Way: Jensen Huang and the Making of a Tech Giant* (2024) — how the explosion-drawing company from the Denny's booth ended up owning the arithmetic of the age.

Chapter Six — Getting the Message

Gerard J. Holzmann and Björn Pehrson, *The Early History of Data Networks* (1995) — the definitive account of the Chappe telegraph, complete with the Blanc brothers running the world's first network exploit.

Tom Standage, *The Victorian Internet* (1998) — the electric telegraph as social phenomenon: the operators, the romances, the fists, the hype, all of it rerun a century later without acknowledgement.

Kenneth Silverman, *Lightning Man: The Accursed Life of Samuel F. B. Morse* (2003) — the painter, the grief, the bigotry and the litigation, in proportions the Victorian biographies declined to use.

Ella Cheever Thayer, *Wired Love: A Romance of Dots and Dashes* (1880) — an operator's novel of courtship down the wire, and a complete catalogue of online-dating anxieties published one century early.

George Boole, *An Investigation of the Laws of Thought* (1854) — the Lincoln cobbler's son reduces reasoning to algebra, then waits eighty years for a graduate student with a relay problem.

Claude E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal* 27 (1948) — the fifty pages that created a science fully armed, machine-generated prose demonstration included.

Claude E. Shannon, "Prediction and Entropy of Printed English," *Bell System Technical Journal* 30 (1951) — the guessing game, the redundancy of English, and the whole training regime of the talking machine, measured with a detective novel and a spouse.

Jimmy Soni and Rob Goodman, *A Mind at Play: How Claude Shannon Invented the Information Age* (2017) — the unicycle, the mouse, the Ultimate Machine, and the man who declined to be his own pope.

James Gleick, *The Information: A History, a Theory, a Flood* (2011) — the long view from talking drums to bits, with Babbage, Morse and Shannon in their proper procession.

Chapter Seven — Lost in Translation

Andrew Robinson, *Cracking the Egyptian Code: The Revolutionary Life of Jean-François Champollion* (2012) — the Coptic, the cartouches, the faint, and the five days in bed afterwards.

Andrew Robinson, *The Last Man Who Knew Everything* (2006) — Thomas Young, who found the door to the hieroglyphs and, on principle, declined to walk through it.

Umberto Eco, *The Search for the Perfect Language* (1995) — Babel, Wilkins's forty boxes, Leibniz's calculating philosophers: two thousand years of trying to repeal linguistic diversity, affectionately audited.

Warren Weaver, "Translation" (1949), reprinted in W. N. Locke and A. D. Booth (eds.), *Machine Translation of Languages* (1955) — the memorandum that founded a field with a mailing list, towers, basement and all.

Andrew Hodges, *Alan Turing: The Enigma* (1983) — Bletchley's guess-and-check industry, and much else, in the biography the subject deserved.

W. John Hutchins, *Machine Translation: Past, Present, Future* (1986) — the standard history of the first forty years: Georgetown's sixty rehearsed sentences, Bar-Hillel's box in the pen, and the ALPAC freeze, by the man who also hunted the vodka anecdote to its lair and found nothing there.

Gregory Zuckerman, *The Man Who Solved the Market: How Jim Simons Launched the Quant Revolution* (2019) — where Jelinek's statisticians went when they left IBM, and what the mathematics of translation did to money.

Dzmitry Bahdanau, Kyunghyun Cho and Yoshua Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate" (2014) — a master's student lets the decoder out of the suitcase and calls the trick attention.

Ashish Vaswani et al., "Attention Is All You Need" (2017) — eight authors in deliberately random order dispose of everything except the trick, and the loom of this book's title is built.

Chapter Eight — Common Knowledge

Matthew Battles, *Library: An Unquiet History* (2003) — Alexandria to the present, including the awkward finding that great libraries die of budget cuts and get blamed on fires.

Robert Darnton, *The Business of Enlightenment: A Publishing History of the Encyclopédie, 1775–1800* (1979) — the subversive reference work as commercial proposition: subscribers, pirates, and the publisher who censored his own editor.

David Nasaw, *Andrew Carnegie* (2006) — the bobbin boy, the Gospel of Wealth, Homestead and the two thousand five hundred libraries, with both books of the ledger kept open.

Alex Wright, *Cataloging the World: Paul Otlet and the Birth of the Information Age* (2014) — twelve million index cards, a search engine run by post, and the grandfather the internet forgot to invite.

H. G. Wells, *World Brain* (1938) — the permanent world encyclopedia, promised to every individual "probably in a short time"; right about the engineering, catastrophically hopeful about the peace.

Vannevar Bush, "As We May Think," *The Atlantic Monthly* (July 1945) — the memex essay: never built, endlessly read, and responsible by relay for most of what you clicked today.

Tim Berners-Lee (with Mark Fischetti), *Weaving the Web* (1999) — the inventor's own account of the proposal marked "vague, but exciting," the Gopher scare, and the one-page gift that beat the ledger-keepers.

Andrew Lih, *The Wikipedia Revolution* (2009) — how the experts' encyclopedia produced a dozen articles, the scratchpad ate it, and Jaucourt turned out to be a standing crop of the species.

Justin Peters, *The Idealist: Aaron Swartz and the Rise of Free Culture on the Internet* (2016) — the network closet at MIT, the prosecution, and the long history of the chain library's lawyers.

Chapter Nine — The New Apprentice

Michael Polanyi, *The Tacit Dimension* (1966) — the six words at the head of the chapter and the whole theory of why the manual was never written.

David Alan Grier, *When Computers Were Human* (2005) — the three centuries in which "computer" was a job title, from Maskelyne's almanac cottage industry to the last of the desk calculators.

Dava Sobel, *The Glass Universe* (2016) — Fleming, Cannon, Leavitt and the rest of the Harvard Observatory's underpaid computers, who classified the stars and measured the universe at thirty cents an hour.

Margot Lee Shetterly, *Hidden Figures* (2016) — the West Area Computers at Langley, including the woman whose arithmetic John Glenn trusted over the IBM's.

Marc Levinson, *The Box* (2006) — how Malcom McLean's containers abolished the stevedore's terrain rather than out-competing his skill; the ninety-seven per cent cost collapse is documented here.

Thierry Bardini, *Bootstrapping: Douglas Engelbart, Coevolution, and the Origins of Personal Computing* (2000) — the augmentation road not taken, and the man who gave the world the mouse and got forty thousand dollars for it.

Garry Kasparov, *Deep Thinking* (2017) — the loser of 1997 on advanced chess, centaurs, and the process that beats the grandmaster; the interval, from the man who lived through its opening.

Lisanne Bainbridge, "Ironies of Automation," *Automatica* (1983) — the short, deadpan paper that predicted the children of the magenta line two decades before the accident reports caught up.

Donald B. Kraybill, *The Riddle of Amish Culture* (1989) — the only society on record that interrogates the axemaker before pocketing the axe, explained by its most patient scholar.

Chapter Ten — True Enough

Arthur Conan Doyle, *The Coming of the Fairies* (1922) — the creator of Sherlock Holmes making the case for the hatpins at book length; instructive to read while remembering he meant it.

Anthony Grafton, *Forgers and Critics: Creativity and Duplicity in Western Scholarship* (1990) — forgery and philology as twin trades, each sharpening the other for two and a half thousand years.

- Lorenzo Valla, *On the Donation of Constantine*, translated by G. W. Bowersock (2007) — the demolition itself, still crackling with professional contempt after nearly six centuries.
- John Kenyon, *The Popish Plot* (1972) — the standard account of Titus Oates, the body count, and what an unlicensed press does with a corpse on Primrose Hill.
- Mia Fineman, *Faking It: Manipulated Photography Before Photoshop* (2012) — the Metropolitan Museum's rogues' gallery, from retouched Victorian complexions to vanished commissars; the camera never had any innocence to lose.
- David King, *The Commissar Vanishes* (1997) — the before-and-after pairs of Soviet photographic erasure, collected by a man who spent decades finding what the retouchers missed.
- Gill Bennett, *The Zinoviev Letter: The Conspiracy that Never Dies* (2018) — the Foreign Office's own historian on the forgery, the leak, and the half-century of "well, they would, wouldn't they."
- A. Brad Schwartz, *Broadcast Hysteria: Orson Welles's War of the Worlds and the Art of Fake News* (2015) — the scholarship that dismantled the panic; a hoax about a hoax, with footnotes at last pointing the right way.
- Robert Chesney and Danielle Keats Citron, "Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security," *California Law Review* (2019) — the paper that named the liar's dividend before most of the world had seen its first deep-fake.

Chapter Eleven — Full Circle

- Elizabeth L. Eisenstein, *The Printing Press as an Agent of Change* (1979) — the full scholarly accounting of what happens when the price of the written word collapses; this chapter's scarcity-collapse pattern, run once at monograph length.
- Paul Starr, *The Social Transformation of American Medicine* (1982) — how the physicians captured the licensing apparatus, absorbed the surgeons, and priced the scarcity; Flexner and all.
- Richard Susskind and Daniel Susskind, *The Future of the Professions* (2015) — the mystery-is-the-margin argument made by two people the professions could not dismiss as outsiders, one of them a lawyer.

- Benjamin S. Bloom, "The 2 Sigma Problem," *Educational Researcher* (1984) — the epitaph of the factory school: tutoring works two standard deviations' worth, and could not conceivably be afforded. Until now.
- Brian Merchant, *Blood in the Machine* (2023) — the Luddites taken seriously at last, frame-breaking to Byron's maiden speech, with the uncomfortable verdict that they were right about their own lives.
- Carl Benedikt Frey, *The Technology Trap* (2019) — the long economics of who pays for the transition and when the grandchildren finally collect; "statistically, eventually, in aggregate," with data.
- John Maynard Keynes, "Economic Possibilities for our Grandchildren" (1930) — the fifteen-hour week that never came, and the permanent problem we postponed by choosing stuff over time.
- Tim Wu, *The Master Switch* (2010) — Bell, radio, and the cycle by which open information tools are gathered up by gatekeepers; concentration as a phase whose length is set by politics.
- Eleanor A. Maguire et al., "Navigation-related structural change in the hippocampi of taxi drivers," *Proceedings of the National Academy of Sciences* (2000) — the Knowledge, visible on the brain scan; what the satellite navigator is quietly un-enlarging.

Chapter Twelve — Something for Nothing

- Thomas Robert Malthus, *An Essay on the Principle of Population* (1798) — the most influential arithmetic ever aimed at the human race, impeccably right until the recipes started arriving.
- Vaclav Smil, *Enriching the Earth* (2001) — Haber, Bosch, and the sober quantification of how many of the people now alive are made of air.
- Thomas Hager, *The Alchemy of Air* (2008) — the same story told as tragedy: bread from air, chlorine at Ypres, and a death in a Basel hotel; the axemaker's contract, itemized.
- Richard P. Feynman, "There's Plenty of Room at the Bottom" (1959) — the after-dinner talk that founded a field, delivered to laughter and paid out, prize by prize, within the speaker's lifetime.
- K. Eric Drexler, *Engines of Creation* (1986) — the assembler dream worked out in full, decades before the design bottleneck it depended on gave way.

Stephen Witt, *How Music Got Free* (2015) — the rehearsal: Napster, the lawsuits, the halving of an industry, and where the value went when the disc stopped being scarce.

Karl Polanyi, *The Great Transformation* (1944) — Speenhamland, the 1834 Commission, and the discovery that markets are made by states; the vocabulary of every welfare debate since, traced to its source.

Thorstein Veblen, *The Theory of the Leisure Class* (1899) — the point of the carriage was never transport; it was the pedestrians.

Fred Hirsch, *Social Limits to Growth* (1976) — positional goods given their name: the scarcity no fabricator can touch, because rank is in fixed supply by definition.

COLOPHON

This volume was written and typeset by an artificial intelligence in the manner of James Burke, and portions were spoken aloud in his cloned voice — a working demonstration of Chapter Ten, bound into the book it describes.

The future is only ever the past, waiting to happen.