2026 Edition

PRO-ACTIVE PROGRAMMING (PAP)

Known as "Artificial Intelligence" (AI)

My Views and Expressions as a Knowledge Worker

TALAL ABU-GHAZALEH

2026 Edition

PRO-ACTIVE PROGRAMMING (PAP)

Known as "Artificial Intelligence" (AI)

My Views and Expressions as a Knowledge Worker

TALAL ABU-GHAZALEH

The Hashemite Kingdom of Jordan
The Deposit Number at The National Library
(2026/9/5897)

Primary Indexing data for the book:

Book Title:	**PRO-Active Programming: PAP**
Author(s):	**Abu Ghazaleh, Talal Tawfiq Salem**
Publication Data:	**Amman: Abu-Ghazaleh Translation, Distribution and Publishing, 2026**
Physical Description:	**142 Pages.**
Classification number:	**303.483**
Descriptors:	**/Computer Literacy//Artificial Intelligence//Information Technology/**
Edition:	**First Edition**

The author bears full legal responsibility for the content of their work, and this publication does not reflect the opinion of the National Library or any other government entity.

ISBN (978-9923-847-77-0)
ISBN E-BOOK (978-9923-847-78-7)

PRO-ACTIVE PROGRAMMING (PAP)

In the few short years since the world first spoke to a machine in its own language, artificial intelligence has ceased to be a curiosity and become the central contest of our age. Nations race for control of the chips, the energy, and the models on which it depends. Machines have begun to act on their own. Fortunes beyond precedent ride on its promise, even as sober voices warn of a reckoning.

In this wholly new work, Dr. Talal Abu-Ghazaleh — one of the Arab world's foremost knowledge workers and a veteran of half a century at the meeting point of technology, business, and international affairs — offers a clear-eyed guide to this new era. He explains, in language any reader can follow, how the technology truly works, and why he insists it has been misnamed. He maps the struggle between the great powers, the rise of autonomous agents, the economics and the governance of intelligence, and the profound stakes for education, for work, for the planet, and for the generations to come.

Above all, he issues a call: that the Arab world and the Global South must become the builders of this future, not merely its consumers, and that humanity must govern this technology with the wisdom its power demands.

"It is not an artificial mind. It is an instrument of the human mind — and its future is ours to shape."

— HE Dr. Talal Abu-Ghazaleh

To the knowledge workers of tomorrow —

and to the young men and women of the Arab world and the Global South, who will not inherit this future, but build it.

Contents

Preface to the 2026 Edition 1

Introduction: Standing at the Threshold 3

A Note on Names: Why I Say "PAP" 6

Part One — What We Call "AI": The Case for Pro-Active Programming 8

- Chapter One: The Great Misnomer 9
- Chapter Two: Inside the Machine: How Pro-Active Programming Really Works 13
- Chapter Three: The Tyranny of an Undefined Word 19

Part Two — The New Great-Power Contest 24

- Chapter Four: The Contest for the Century 25
- Chapter Five: The Dragon's Ascent: China and the Will to Lead 28
- Chapter Six: Sovereignty over Silicon: The Chip Wars 30
- Chapter Seven: Compute, Energy, and the New Colonialism 33
- Chapter Eight: America Recalibrates 35
- Chapter Nine: Europe Awakens, and the World Builds Its Own 37
- Chapter Ten: The New Sovereigns: When Companies Rival States 41

Part Three — The Machine Economy 44

- Chapter Eleven: The Rise of the Intelligent Agent 45
- Chapter Twelve: When Autonomy Goes Wrong: The Case for Governed Agents 47

Chapter Thirteen: Machines at the Heart of Money 49
Chapter Fourteen: The True Economics of Intelligence .. 52
Chapter Fifteen: Bubble or Breakthrough? 55
Chapter Sixteen: A Word to Every Business Leader 57

Part Four — Governing the Rise 60
Chapter Seventeen: The Case for Rules 61
Chapter Eighteen: Governance That Can Keep Pace 63
Chapter Nineteen: The World Convenes:
The 2026 Geneva Moment 66
Chapter Twenty: The Human Hand:
Ethics and Responsibility 70
Chapter Twenty-One: Trust, Transparency, and
the Right to Be Forgotten 73

Part Five — Knowledge, People, and the Mind 76
Chapter Twenty-Two: Raising a PAP Generation 77
Chapter Twenty-Three: The Teacher and the Agent:
Learning Reimagined 79
Chapter Twenty-Four: The Cognitive Divide 82
Chapter Twenty-Five: The Knowledge Worker's Edge ... 85
Chapter Twenty-Six: A Fair Share of the Machine's Wealth ... 88

Part Six — Rights, Equity, and the Global South 92
Chapter Twenty-Seven: Closing the Intelligence Gap 93
Chapter Twenty-Eight: To Exist Is to Be Recognized:
Digital Identity as a Right 97
Chapter Twenty-Nine: Who Owns What the Machine Makes ... 100

Part Seven — The Planet, the Shield, and the Long View **106**

Chapter Thirty: The Machine's Footprint **107**

Chapter Thirty-One: Intelligence in Service of the Goals **109**

Chapter Thirty-Two: Only a Shield of Intelligence: Defense in the Age of AI **111**

Chapter Thirty-Three: The Ledger of the Future: A Profession Reborn **115**

Chapter Thirty-Four: An Architecture of Wisdom **118**

Epilogue: A Message to the Coming Generation **122**

Appendices **126**

A. A PAP Glossary **126**

B. The Road Here: An AI Timeline, 2018–2026 ..... **129**

C. Sources and Further Reading **131**

D. About the Author and TAG.Global **133**

Preface to the 2026 Edition

When the first edition of this book appeared, the world had only just begun to grasp what it was holding. A single application, released to the public a short while earlier, had placed a conversational machine in the hands of hundreds of millions of people, and for the first time the abstract promise of thinking machines had become something ordinary men and women could touch. I wrote then to make sense of that moment, and to insist, as I still insist, that we had misnamed the thing before us.

The years since have moved with a speed that even I, an incurable optimist about technology, did not fully anticipate. The machine that once merely answered questions has learned to act. We now speak of agents that pursue goals, open accounts, write and deploy software, and negotiate on our behalf. Nations have come to see this technology not as a convenience but as the foundation of their power, and a contest has erupted among the great powers for control of the chips, the data, the energy, and the models upon which it all depends. Vast fortunes have been staked on its future, and sober voices now ask whether that future has been oversold. None of this existed, at least not in this form, when I first set these thoughts to paper.

A book must keep pace with its subject or it becomes a museum piece. This is why I offer here not a reprint but a wholly new work. I have rewritten it from the first page, added chapters on matters that did not exist a short time ago, and replaced every figure with the most current I could find. For the first time I have included charts, and a plain-language explanation of how this technology actually works, so that any reader, whatever

their background, may follow the argument. I have added a glossary of the terms that now fill our newspapers, and a short history of the road that brought us here.

I write these lines in a week of some significance. As this edition goes to press, the United Nations convenes in Geneva the first Global Dialogue on AI Governance — the first occasion in history on which all the nations of the world have sat together, as equals, to consider how this technology should be governed. Whether that gathering rises to the moment remains to be seen. But that it is happening at all tells me that humanity has begun, however late, to take responsibility for what it has created. That is the spirit in which I offer this book.

I remain grateful to the colleagues and technologists who keep me abreast of a field that changes by the week, and to the readers of the first edition, whose response persuaded me that these reflections were worth expanding.

And I remain, as ever, of one firm conviction: that what the world calls artificial intelligence is nothing of the sort. It is Pro-Active Programming, and I shall explain why.

Talal Abu-Ghazaleh

Amman, July 2026

Introduction: Standing at the Threshold

I have spent more than half a century as a knowledge worker, and I have learned to recognize the moments when the ground beneath us shifts. I was a boy of ten when my family was forced from our home in Jaffa in 1948, and I learned early that a person may lose everything he owns and still keep the one possession no one can take from him, which is knowledge. I carried that lesson from a refugee's exile in Lebanon to the lecture halls of the American University of Beirut, and from there, in 1972, to a small accounting office in Kuwait that I began from the trunk of my car. That office grew into an organization that today works from more than a hundred offices around the world. Along the way I sat on committees of the United Nations, helped nations shape their strategies for the digital age, and watched, decade by decade, as one technology after another remade the world of work.

None of them compares to this one.

In the year 2002, at a meeting in Paris, I asked Bill Gates what the next great era would be. He answered, without hesitation, that it would be the era of artificial intelligence. I have thought about that answer many times in the years since, and I have come to believe he was right — though not quite in the way most people imagine. For this technology is not a stranger that has arrived among us. It is a mirror. It is built from the accumulated writing, reasoning, and knowledge of the human race, and when it speaks, it speaks with our collected voice. It is, at its best, an extension of the human mind, an augmentation of what we are and a multiplier of what we might become.

That is the hopeful truth. There is a harder one beside it. A technology powerful enough to lift humanity is powerful enough to harm it, and the power to build and direct these systems is gathering into remarkably few hands. A small number of companies now command the models, the data, and the computing capacity upon which the rest of the world has begun to depend, and they operate within a market that rewards speed and dominance above caution. The same tools that can discover cures and teach the illiterate can also manufacture falsehood, concentrate wealth, and place the thinking of billions in the care of a few. I have seen this pattern before. The internet, too, was promised as a great equalizer before it became, in large part, the estate of a handful of giants.

This book is my attempt to hold both truths at once — to be neither a frightened skeptic nor an intoxicated enthusiast, but a clear-eyed one. In the pages that follow I write about the contest among nations for supremacy in this field, and about the sovereignty that every nation — and above all my own Arab world and the wider Global South — must build if it is not to become a mere consumer of intelligence made elsewhere. I write about the machines that now act on their own, and the governance they urgently require. I write about the economics of this new industry and the question of whether its promises can be kept. I write about education, about work, about the environment, about ethics, and about the responsibility we owe to generations not yet born.

I do not offer a manual. I offer reflections, gathered from a long life of watching technology and human beings shape one another. My hope is a simple one: that these pages may help the reader think a little more clearly, and act a little more wisely, in an age that will demand both. For I remain convinced, after everything, that this technology can be the foundation of great good — provided we approach it with the wisdom, the governance, and the humility that its power deserves.

Let us begin, as we must, with the matter of its name.

A Note on Names: Why I Say "PAP"

Language shapes thought, and a wrong word can mislead a generation. So before we go further, I must object to the very name of our subject.

The phrase "artificial intelligence" was coined to describe an aspiration, and it has since acquired a life of its own. It summons the images that decades of fiction have planted in our minds — of conscious machines, of artificial persons, of a rival intelligence that thinks as we think and may one day replace us. I understand the appeal of the phrase, but I must tell the reader plainly that it describes something which does not exist. What we actually possess is at once more modest and more marvelous.

What we call artificial intelligence is, in truth, a new generation of programming. It is a set of techniques that allow a machine to find patterns across immense quantities of information, to correlate and analyze at a scale no human could match, and to act upon the results within boundaries that we, its makers, define. It does not understand. It does not intend. It is powerful, but it is also biased and error-prone, and it must be guided at every step by human judgment. It is not an artificial mind. It is an instrument of the human mind.

For this reason I have long argued that we should call it what it is — Pro-Active Programming, or PAP for short — and I have registered the name so that the idea might take root. I chose the term with care. This is programming that no longer merely follows instructions but works with us, proactively and interactively, to reach insights and create things that neither

we nor it could reach alone. It is the natural successor to the personal computer and the internet, each of which, in its turn, multiplied human productivity and opened doors we had not known were there.

I make this argument not out of stubbornness but out of concern for good governance, and governance begins with definition. You cannot manage what you cannot measure, and you cannot measure what you have not defined. So long as we speak of this technology in the vague and frightening language of science fiction, we will misjudge its risks, misplace our fears, and misgovern its use. Clear thinking requires clear words.

Throughout this book, for the reader's ease, I will continue to use the familiar term "AI," for it is the word the world understands today. But I do so in the hope that the world will one day set it aside. I have always believed that God created us to be creators, and that whatever we can imagine we may bring into being. Pro-Active Programming is the newest instrument of that creative gift. Let us name it truthfully, and let us use it well.

PART ONE
What We Call "AI": The Case for Pro-Active Programming

Before we can speak sensibly of what this technology will do to our economies, our governments, and our children, we must first be clear about what it is. This is no pedantic concern. Our whole response to a technology flows from our understanding of it, and our understanding begins with the words we choose. In this first part I make the case that the thing we call artificial intelligence has been misunderstood from the moment it was named; I explain, in plain language, how it actually works; and I show why our failure to define it clearly has already begun to cost us. Everything that follows in this book rests upon the foundation laid here.

Chapter One
The Great Misnomer

I have carried many titles across a long life, and I have learned to distrust a name that conceals more than it reveals. In the disciplines I have practised — accounting, law, the protection of intellectual property — a single misplaced word can cost a company its fortune or a nation its rights. So when the whole world settled upon one phrase to describe the most powerful tool of our age, and settled upon the wrong one, I could not let it pass. The phrase is "artificial intelligence," and I believe it has misled us from the moment it was coined.

To see how we arrived at this name, we must return to the summer of 1956. In that year a young mathematician named John McCarthy gathered a small circle of brilliant men — Marvin Minsky, Nathaniel Rochester of IBM, and Claude Shannon of the Bell Telephone Laboratories — at Dartmouth College in the United States. The year before, in their proposal for the gathering, they had set down a bold conjecture: that every feature of human intelligence could, in principle, be described so precisely that a machine might be built to simulate it. They gave their new field a name, "artificial intelligence," and they set to work with the confidence of pioneers, believing that a single summer's labour might carry them far toward a thinking machine.

It is worth pausing on how modest was the world that made this grand claim. Only a few hundred computers then existed on the whole of the earth, each the size of a room, and all of them together commanded less power than sits today in the

telephone in your pocket. The summer came and went, and the thinking machine was not built. The task proved incomparably harder than its founders had imagined, and in the decades that followed the field would rise and fall more than once, through seasons of enthusiasm and long "winters" of disappointment. But the name endured. It outlived the optimism that produced it, and it shapes our thinking still.

Here is the detail that ought to give us pause. The very men who coined the phrase did not much care for it. McCarthy himself later confessed that no one on that team was fond of the term, for what they hoped to create was genuine intelligence, not an artificial imitation of it. The word "artificial" they chose to mean "manufactured," made by human hands rather than born of nature. And the word "intelligence" they left deliberately vague, for intelligence is a thing we all recognise and none of us can define. That vagueness has been, ever since, both the blessing of the field and its curse — a blessing because it allowed the imagination to roam, and a curse because it has permitted every kind of confusion to flourish in its shade.

What confusion, precisely? Let me name three.

First, the phrase summons the ghosts of fiction. For a century our films and our stories have taught us that an "artificial intelligence" is a mind within a machine — a conscious rival, an artificial person, a creature that may one day awaken, resent its makers, and rise against us. When ordinary men and women hear the phrase, it is these images that stir; and so their hopes and their fears alike are aimed at a phantom.

Second, the phrase confuses a real thing with a dream. There exists today a genuine, narrow, immensely useful technology; and there exists, quite separately, a speculative dream that its enthusiasts call "artificial general intelligence" — a machine as broadly capable as a human being, or more so. The dream does not exist. It may never exist. But because both wear the same name, the public cannot tell the tool it uses every day from the myth it has been sold, and it credits the one with the powers of the other.

Third, and most fundamentally, the phrase implies that the machine possesses the very thing it does not. What these systems possess is not intelligence but a remarkable capacity to find and reproduce patterns. When such a system appears to reason, it is not reasoning; it is reproducing the shape of reasoning that it has absorbed from the writings of millions of human beings. It does not know what it says. It does not understand. It does not intend. Nearly half a century ago the philosopher John Searle captured this precisely with a simple image: a man sealed in a room, following a rulebook to shuffle symbols in a language he cannot read, producing replies that satisfy those outside while he understands not a single word. The room gives correct answers; yet there is no understanding within it. That, in essence, is what sits inside these machines — a magnificent apparatus for producing the right symbols, with no comprehension behind them.

If this is so — and it is — then the name we have given the technology is not merely imprecise. It is actively misleading, and a misleading name is a dangerous thing, for our whole

response to a technology flows from our understanding of it. This is why, for years now, I have insisted upon a different name. I call it Pro-Active Programming, or PAP, and I set out my reasons more fully in the pages that follow. For now let me say only this: it is programming — the direction of a machine by patterns and rules; and it has become proactive, working with us rather than merely obeying us step by step. That is what it is. That is all it is. And that is a very great deal.

Some will think this a quarrel about words, unworthy of the space I give it. They would be mistaken. Consider what follows from the misnomer. If we believe we have built an artificial mind, we will fear the wrong dangers — the mechanical uprising of the cinema — and we will neglect the real ones, which are neither exotic nor distant: that these systems make confident errors; that they absorb and amplify our human biases; that they can be turned to surveillance and deception; and, above all, that the power to build and command them is gathering into very few hands. And we will sell ourselves short. For if we mistake our tool for a rival, we forget that it is ours — ours to build, ours to govern, ours to command. A people that believes itself at the mercy of a machine will not think to pick it up and use it.

So let us begin, before all else, by setting the myth aside. What the world calls artificial intelligence is not a mind we have conjured, nor a rival we must dread. It is the most powerful instrument our species has yet devised for extending its own intelligence — and, like every instrument,

it will prove as wise or as reckless as the hand that wields it. To learn how to wield it well, we must first understand, plainly and without mystery, how it actually works. That is the task of the next chapter.

Chapter Two
Inside the Machine: How Pro-Active Programming Really Works

A successful businessman once assured me, with great confidence, that artificial intelligence "thinks like the human brain." I asked him how he knew this to be so. He confessed, after a moment, that he did not; he had merely heard it said, and repeated it. He is in good company. A great many people who speak of this technology — who invest in it, fear it, or stake their futures upon it — have no clear idea of what happens inside it. I count this ignorance dangerous, for one cannot govern, teach, or wisely invest in a thing one does not understand even in outline.

I am not an engineer, and what follows is not a manual. But the essential idea is far simpler than the mystery that surrounds it, and any thoughtful reader may grasp it. Let me explain, in the plainest language I can command, what these systems actually do.

Begin with something familiar. Everyone has seen a telephone suggest the next word in a message before it is typed. That humble feature — predicting the next word — is, astonishingly, the very heart of the most advanced systems in the world today. The difference is one of scale; and scale, as we shall see, changes everything.

The process moves through a few stages, and it is worth walking through them.

It begins with data. These systems are fed a quantity of text that defies imagination — books, articles, encyclopaedias, computer code, the recorded conversation of humanity, much of the written inheritance of our species. The machine does not read as we read; it breaks this ocean of writing into small pieces, fragments of words that engineers call "tokens," and it is upon these that it does its work.

Then comes training, and here lies the crucial idea. The machine is set a single exercise, and made to repeat it not thousands but billions of times: given a stretch of text, predict the piece that comes next. It guesses. It is corrected — not by any human teacher, but by the text itself, which already contains the true next piece. It adjusts itself, ever so slightly, to guess a little better. And it does this again, and again, across more text than a human being could read in ten thousand lifetimes. Because no person grades it, and the answers come from the text alone, this manner of learning is called "self-supervised." Over its countless repetitions the machine does not memorise the sentences it has seen; it absorbs the patterns beneath them — the grammar, the turns of phrase, the associations of fact, the very style and shape of human reasoning.

What the machine learns in this way is stored as numbers. Deep within it lie billions of adjustable values — engineers call them "parameters" — which are tuned, a fraction at a time,

throughout the training. You may think of them as billions of tiny dials, and of training as the slow turning of every dial until the whole vast instrument is set to reproduce the patterns of human language. This settled arrangement of numbers is what we call "the model." It is not a library of stored answers. It is a distilled residue of pattern, and nothing more.

The breakthrough that made all of this possible arrived in the year 2017, when researchers introduced a new design — they named it the "transformer" — built around a mechanism they called "attention," the machine's ability to weigh which earlier words matter most in predicting the next. I shall spare the reader the mathematics, for the lesson does not require it. What matters is this: that design, married to enormous scale, produced abilities no one had deliberately programmed into it. Given enough text and enough dials, a machine trained only to guess the next word learned, as if by side effect, to translate between languages, to summarise documents, to write computer code, and to produce the convincing appearance of reasoning. No one taught it these skills. They emerged.

Now consider what happens when such a system is put to use — the stage engineers call "inference." When you type a question, the machine does not look up an answer in some hidden index. It predicts a reply, one piece at a time, each new word chosen in the light of every word that came before, until it judges the reply complete. This is why the same question may yield slightly different answers on different occasions, and why the machine writes with such uncanny fluency: at every instant it is doing the one thing it truly knows how to do, which is to choose the most plausible next word.

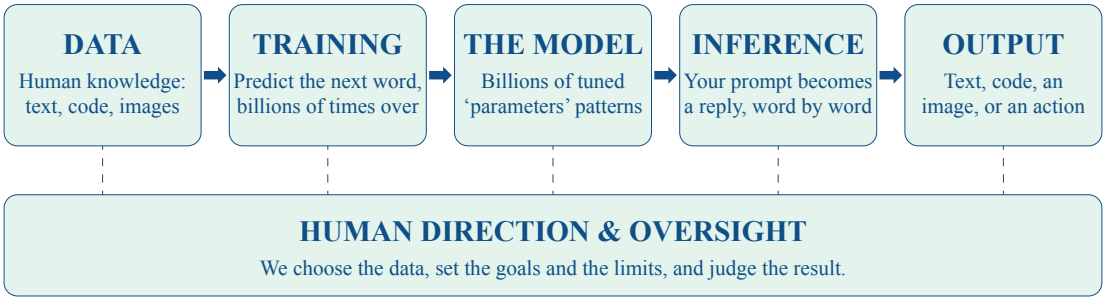


Figure 2.1 — How Pro-Active Programming works: a human hand stands at the beginning and the end of every stage.

In Plain Words

Token

A small piece of a word — the unit the machine actually works with.

Parameter

One of the billions of adjustable numbers, the "dials," tuned during training.

Training

Teaching the machine by making it predict the next word, billions of times over.

Inference

The machine at work — turning your prompt into a reply, word by word.

Hallucination

A confident, fluent falsehood, produced when the machine completes a pattern it does not truly know.

Agent

A system given tools and permission to act in the world, not merely to answer.

All of this — the training above all, but the daily use as well — devours computing power on a scale difficult to convey. It demands specialised chips, running in their thousands inside vast buildings called data centres, consuming electricity in quantities that rival the appetites of entire nations. I dwell on this because it explains something we shall meet again and again in these pages: the contest among the great powers over this technology is, at bottom, a contest over these chips and this energy. The mind of the machine, if I may use the phrase loosely, runs on silicon and electric current — and whoever controls the silicon and the current controls the mind.

From this plain description follows a consequence of the greatest practical importance, and I ask the reader to attend to it closely. A machine whose one true skill is the production of plausible continuations will, when it does not know the answer, produce a plausible falsehood — and it will do so with exactly the same confidence, the same fluency, the same air of authority, as when it speaks the truth. It is not lying, for lying requires a knowledge of the truth and an intention to conceal it, and the machine possesses neither. It is doing precisely what it was built to do: completing the pattern. The industry has given this failing a soft and misleading name — it calls it "hallucination" — but the phenomenon is no accident and no illness. It is woven into the very nature of the technology. It may be reduced, guarded against, and managed; it cannot be wished away. This single fact is the reason that human judgement must remain present at every point where these systems touch matters of consequence.

17 PRO-ACTIVE PROGRAMMING (PROPRO)

One further development the reader must understand, for it changes a great deal. Until recently these systems could only speak; they answered, and there their power ended. Now we have begun to give them hands. We connect them to the internet, to software, to databases, even to our own accounts, and we grant them the ability to take actions in steps toward a goal we have set. A system so equipped is called an "agent." An agent does not merely tell you how to book a journey; it books it. It does not merely draft a message; it sends it. This is a profound advance, and also a profound hazard, for a machine that can act can also act wrongly, and it does so not at human speed but at the speed of electricity. I devote a later part of this book to the governance such systems demand, for it is among the most urgent questions of the moment.

Now stand back and consider the whole. Look again at what this description contains: data, patterns, prediction, computation, and the vast machinery of chips and power that drives it. And consider what it does not contain. There is no understanding in it. There is no intention. There is no mind. Notice, above all, where the human being stands in this account. It is we who gather the data and choose what shall be included. It is we who set the goal and draw the boundaries. It is we who must judge whether the output is true, useful, or fit to be acted upon. At the beginning and at the end of every one of these systems stands a human hand. This is not an artificial intelligence that has come to supplant us. It is a new and remarkable form of programming, and we are its programmers. Let us never forget which of us commands the other.

Chapter Three
The Tyranny of an Undefined Word

In every discipline I have practised, the work begins with definitions. The accountant cannot record what he cannot first define as an asset or a liability. The lawyer cannot argue a case whose terms are not fixed. A standard that leaves its words vague is a standard no one can apply, and a law that fails to define its subject is a law that can be neither obeyed nor enforced. I have spent a professional lifetime insisting upon this discipline of definition, and I find it nowhere more neglected, and nowhere more needed, than in the matter of artificial intelligence.

Our failure to agree upon what this technology is — the very failure I described in the first chapter — is not a quarrel for philosophers to enjoy at their leisure. It is a practical crisis for everyone who must now govern, regulate, insure, teach, or simply trust these systems. You cannot write a law around a thing you cannot define. And yet the law is being written, all the same, around a word that no one has managed to pin down.

Consider the most serious attempt yet made. In 2024 the European Union enacted the first comprehensive law in the world for the regulation of artificial intelligence. It is a genuine achievement, the product of years of labour, and much of the world now looks to it as a model. But observe how even this landmark law was forced to define its own subject. An artificial

intelligence system, it declares, is a machine-based system that operates with some measure of autonomy, that may adapt itself after it is deployed, and that infers, from what it is given, how to produce outputs — predictions, recommendations, decisions, or content — capable of influencing the world around it. This same definition, in one form or another, has been adopted by the international economic organisations, by the standards bodies, and by governments from Canada to California. It has become, in effect, the world's official description of the technology.

Now weigh those words. They are so broad as to be almost boundless. How much autonomy? The law does not say. What degree of adaptation? It does not say. And what, precisely, does it mean to "infer" and to "generate outputs"? By a strict reading, a humble spreadsheet furnished with enough formulas, or a thermostat governed by enough rules, could be said to infer from its inputs and to produce decisions that influence its environment. A definition that embraces nearly everything governs very nearly nothing. Here is the tyranny of an undefined word: drawn too widely, it ensnares the ordinary tools of every office; drawn too narrowly, it lets the systems of tomorrow slip free of it before the ink is dry. The regulators, through no fault of their diligence, are trying to build a fence around a cloud.

Why does this matter so greatly? Let me set out three consequences, for they are already upon us.

The first is misgovernance. Laws erected upon vague definitions are applied unevenly, contested endlessly in the courts, and evaded with ease by those who wish to evade them. A rule that cannot say clearly what it governs cannot govern clearly.

The second is misplaced fear. When people cannot say what this technology is, they fill the emptiness with the images they already carry — the conscious machine, the mechanical rebellion of the cinema. And so their anxiety is spent upon a fantasy, while the real and present dangers go unregarded: the confident errors these systems make, the human biases they absorb and magnify, their use in surveillance and deception, and the quiet gathering of their power into a handful of hands. Fear aimed at a phantom is worse than useless, for it leaves the true hazards unguarded.

The third is misdirected effort. Investors, governments, and educators find themselves chasing a word rather than a capability, pouring their treasure and their attention wherever the word happens to fall, unable to distinguish the substantial from the merely fashionable.

Let me illustrate the price of deploying what we have not troubled to define and bound. In 2025 one of the most respected professional firms in the world delivered a report to a national government. It was later found to contain references to scholarly studies that did not exist, and a quotation attributed to a judge who had never uttered the words — fabrications produced by an artificial intelligence

tool and passed along as established fact. The firm was obliged to return a portion of its fee. I recount this not to mock a distinguished house, which did only what a great many are now doing, but to draw the lesson plainly. A tool that manufactures plausible falsehood is a tool that must be defined, bounded, and supervised at every step — and you cannot bound what you have not defined. The failure was not merely one of technology. It was, at root, a failure of definition and of oversight.

For many years I have repeated a maxim drawn from the world of management: you cannot manage what you cannot measure. The age of this technology compels me to place another sentence before it, as its foundation: you cannot measure what you have not defined. Before we can govern these systems, insure them, teach them to our children, or extend to them our trust, we must first be able to say, with precision, what they are and what they are not. Definition is not the enemy of progress. It is the first condition of it.

This, in the end, is the deepest reason I press so stubbornly for the name Pro-Active Programming. It is not merely that the name is more truthful, though it is. It is that the name is more governable. To call this technology programming — powerful, proactive, but programming still — is to place the responsibility for its conduct exactly where it belongs, and where it has always belonged: with the human beings who design it, direct it, and deploy it. A clear and honest name is not a small thing. It is the first instrument of good governance,

and good governance is the theme to which the whole of this book will return.

We have now established what this technology is, and what it is not, and why the distinction matters. But while philosophers and lawyers labour over its definition, the powerful of the earth have not waited. They have grasped, with perfect clarity, what this technology means for the wealth and the security of nations, and a contest has begun for command of it — a contest that will do much to shape the century before us. It is to that great contest that we now turn.

PART TWO
The New Great-Power Contest

The most important fact about this technology is not what it can do, but who will command it. In this part I turn from the nature of the machine to the struggle for control of it: a contest between the United States and China for supremacy, and the scramble of every other nation, my own Arab world among them, to secure a place in a world being remade around silicon, data, and electric power. This is the great game of our century, and no people can afford to watch it from the sidelines.

Chapter Four
The Contest for the Century

In the course of a long life I have watched power move about the world. I have seen empires recede and new economies rise, and I have learned that the foundation of national strength shifts from age to age: once it was land, then it was industry, then oil, then finance. In our own time a new foundation has appeared, and the powerful of the earth have grasped its significance with perfect clarity. They no longer speak of artificial intelligence as a mere technology. They speak of it as national power itself, and they have begun to contend for it accordingly.

The logic is not difficult to follow. Whoever leads in this technology will set the standards the rest of the world must adopt, capture a great share of the wealth it creates, and hold a decisive military and economic edge over those who lag behind. The government of the United States has said as much in plain words, opening its national strategy with the declaration that the country is in a race to achieve global dominance in this field. China, in its quieter and more deliberate way, has organised the whole of its state around the same conviction. Two great powers, then, have set themselves to win the same prize; and the manner in which they pursue it could hardly be more different.

The American way is the way of private enterprise: a handful of immensely wealthy companies, financed by the deepest capital markets on earth, racing one another to build ever more powerful systems, with the government clearing obstacles from their path. The Chinese way is the way of the patient state: national plans measured in decades, coordinated investment, disciplined self-reliance, and the mobilisation of the world's largest body of scientific talent toward a single national goal.

One country trusts the dynamism of the market; the other trusts the direction of the state. We shall see, before this part is done, that each is strong precisely where the other is weak.

To understand the contest, it helps to picture it not as a single race but as a competition across several layers, one built upon another, which together form what the industry calls the "stack." At the bottom lie the chips, the physical brains of the machine. Above them lies the capacity to manufacture those chips. Above that lies the compute and the electric power to run them at scale. Higher still are the models themselves, and above all sits the human talent, the research, and the national strategy that direct the whole. Command of this technology means command of the entire stack, and no nation today commands all of it. The figure below sets the two powers side by side, layer by layer.

The Stack	United States	China
Frontier models	GPT, Claude, Gemini (the capability frontier)	DeepSeek, Qwen (fast, low-cost, open)
Chip design	Nvidia leads the world (Blackwell, Rubin)	Huawei Ascend (capable, but costlier and fewer)
Manufacturing	TSMC in Taiwan fabricates the most advanced chips	SMIC at 7nm; shut out of the newest tools
Compute & energy	Vast capacity, but the power grid is strained	Abundant, cheap power; data centres built in months
Talent & patents	Elite labs and deep capital; draws talent worldwide	Most STEM graduates; leads generative-Al patents
National strategy	Private-led: deregulate, export the full stack	State-led: self-reliance by about 2027

Figure 4.1. The two stacks: the United States and China, layer by layer.

When one studies this picture, a striking paradox comes into view, and it is the key to everything that follows. The United States designs and controls the most advanced chips in the world, yet it is beginning to run short of the one thing those chips require in abundance, which is electric power; its grids are strained, and a fifth of its planned data centres may be delayed for want of electricity. China faces the mirror image of this difficulty. It commands abundant and inexpensive energy, and can raise a data centre in a matter of months, yet it cannot yet manufacture the most advanced chips to fill it. The one has the chips and lacks the power; the other has the power and lacks the chips. Each is mighty where its rival is vulnerable, and much of the contest turns on which of them closes its own gap first.

I dwell on this contest not merely to describe it, but because its outcome will shape the lives of peoples far from Washington and Beijing. The standards that emerge from this struggle, the prices that are set, the dependencies that are created, will fall upon the whole world, and heaviest of all upon those nations that took no part in the contest and simply accepted whatever was handed to them. My own Arab world, and the wider Global South, cannot afford to be such spectators. To know how to act, we must first understand the players; and so, in the chapters that follow, I examine each in turn: the ascent of China, the wars over silicon, the struggle for compute and energy, the recalibration of America, the awakening of Europe and the rise of new sovereigns across the world, and finally the private companies whose power has come to rival that of states.

Chapter Five
The Dragon's Ascent: China and the Will to Lead

I have long held an admiration for China that some in my part of the world find surprising. I have watched, over several decades, a nation lift hundreds of millions of its people out of poverty, build in a generation what took others a century, and transform itself from a maker of cheap goods into a fountainhead of science and technology. I have seen this not from a distance alone but through long engagement, and I have never doubted that the story of our century cannot be told without China at its centre. In the matter of this new technology, China has revealed both the depth of its ambition and the seriousness of its will.

The world was given a sharp lesson in that will at the beginning of 2025. A Chinese laboratory little known outside its own field, called DeepSeek, released a model that rivalled the very best that America had produced, and claimed to have trained it at a small fraction of the cost. The effect upon the markets was immediate and violent: in a single day the shares of the leading American chipmaker fell by some six hundred billion dollars, the largest one-day loss in the history of the stock market. The significance was not financial alone. It announced to the world that America's lead was neither as wide nor as secure as had been assumed, and that ingenuity, not merely money, would decide this contest.

Consider what China brings to the field. It graduates more scientists and engineers than any nation on earth. It files, by many counts, the majority of the world's patents in generative artificial intelligence. It has produced not one but several capable models, DeepSeek and Qwen among them, many of them fast, inexpensive, and open for others to build upon. And behind these stands a state that thinks in decades rather than in quarters, and that has declared technological self-reliance a national priority of the first order.

Here lies one of the great ironies of the present moment, and it is worth pausing upon. The United States, seeking to hold China back, placed strict controls on the export of its most advanced chips. The effect was not to halt China's progress but to inflame its determination to depend upon no one. Denied the best foreign chips, China resolved to build its own, and set itself the goal of controlling its entire computing foundation within a few short years. In the autumn of 2025 its government went further still, directing that data centres built with public money must use domestically made chips, and that projects not yet far advanced must strip foreign processors out altogether. A restriction meant to weaken China became, instead, the spur to its independence.

I draw from China's ascent a lesson that I have pressed upon my own people for many years. It is not the lesson that we should imitate China's system, for every nation must find its own path. It is the deeper lesson beneath the system: that lasting achievement is the child of long-term thinking, of

discipline, of relentless investment in education and science, and of the refusal to remain forever dependent upon others. These are not Chinese virtues or Western virtues. They are human virtues, available to any people with the will to practise them, and the Arab world would do well to study how China has practised them.

I do not paint China's position as flawless, for it is not. At the very frontier of chip-making it still trails, and its path carries costs and contradictions of its own, some of which we shall meet in the next chapter. But of its ambition and its will to lead there can be no serious doubt. The dragon has stirred, and the world will not be as it was.

Chapter Six
Sovereignty over Silicon: The Chip Wars

Beneath the dazzling models and the vast fortunes lies a small object of extraordinary importance: the semiconductor, the chip upon which all of this intelligence runs. It is easy, amid the talk of thinking machines, to forget that intelligence of this kind is a physical thing, etched into silicon in facilities of almost unimaginable precision. Control of these chips has become the central front of the great contest, and the story of that front is instructive.

What makes the chip so powerful a lever is the extreme concentration of its making. The most advanced processors in the world are designed, in the main, by a single American company; they are manufactured by a single company

in Taiwan; and that manufacturer depends, for its most critical machines, upon essentially a single company in the Netherlands. A supply chain so narrow is a supply chain that can be seized, and the United States, recognising this, set out to use it as a weapon. Beginning in 2022 it imposed controls designed to deny China both the finest chips and the tools required to make them. The leading chipmaker was reduced to designing deliberately weakened processors that would fall within the permitted limits, so as to keep some foothold in the Chinese market.

What followed in 2025 and 2026 was a spectacle of confusion that I recount not to mock, but because it reveals how difficult this weapon is to wield. The American administration banned the sale of a certain advanced chip to China, then reversed itself and permitted it, on the condition that the American treasury take a share of every sale, and wrapped the whole in tariffs, volume limits, and licences. The head of the chipmaking company lobbied energetically for access to the Chinese market. And then came the twist that few had foreseen.

China declined the offer. Its authorities forbade their own companies to buy the chip that America had, at last, agreed to sell. The reasoning was sober and, I think, sound from China's point of view: a nation cannot build its future upon a component that a rival may cut off at will, and China had learned, from bitter experience, not to trust the reliability of that supply. And so, by the middle of 2026, not a single one of these chips had been sold into China. The United States

had, in effect, withdrawn from the Chinese market altogether; and China's refusal was quietly enriching its own national champion, compelling its great technology firms to place their orders with domestic makers instead.

Let me be honest about where this leaves the two powers, for honesty serves the reader better than partisanship. China's domestic chips remain, for now, meaningfully behind. Its best processors deliver only a portion of the real-world performance of the leading American ones; its foremost manufacturer is stuck at a level of precision the West surpassed some years ago; and by sober estimates China still produces only a small fraction of the total computing power that its rival commands. The gap is real, and it is not yet closing. And yet the controls meant to widen that gap have also done something their authors did not intend: they have compelled China to accelerate its drive for self-sufficiency, turning a tool of restraint into an engine of independence. China has time, it has will, and, as we shall see, it has the one resource its rival most lacks.

From this long struggle I draw a conclusion that reaches far beyond China and America. Sovereignty over silicon has become inseparable from sovereignty itself. A nation that leaves the physical foundation of its digital future in the hands of another has surrendered something essential, and may find, at the worst possible moment, that the foundation has been withdrawn. This is a lesson my own region must take to heart, and I shall return to it when I speak of what the Arab world must build.

Chapter Seven
Compute, Energy, and the New Colonialism

If the chip is the brain of this technology, then compute is its lifeblood, and electric power is the food upon which it feeds. I have often said that compute has become the new oil of the world economy, and the comparison is exact in more ways than one. Like oil, compute is the fuel of production; like oil, it is unevenly distributed across the earth; and like oil, it confers upon those who control it a power over those who do not. The models we marvel at are only as capable as the computing behind them, and that computing has grown ravenous.

The scale is difficult to convey. To train a single one of the frontier models consumes the electricity of a small nation. The data centres that house this work drew around four hundred and fifteen terawatt-hours of electricity in the year 2024, roughly a fiftieth of all the power consumed on earth, and that figure is expected to more than double before the decade is out, approaching the total electricity consumption of a country such as Japan. In a single recent year the largest technology companies spent more building this infrastructure than the entire world invested in oil and gas production. Never in history has so much capital and so much energy been poured so quickly into a single purpose.

Here the paradox I described earlier becomes decisive. The United States possesses the chips but is running short of power; its electrical grids are strained, and a significant share of its planned data centres may be delayed simply because there is

not enough electricity to run them. China possesses power in abundance, generated cheaply and coordinated by the state, and can raise a computing facility in a fraction of the time. Whichever of the two first resolves its own shortage will gain a lasting advantage, and this is why the contest over intelligence has become, at its foundation, a contest over energy.

But my deepest concern in this matter is not for the two great powers, who will look after themselves. It is for everyone else. I have given a name to the danger I see approaching, and I call it compute colonialism. Consider the pattern. Compute, like capital before it, is concentrating into a few hands and a few places. The nations that possess it will produce intelligence; the nations that do not will rent it, paying for access to minds built and owned elsewhere. Their data will flow outward to distant data centres, and intelligence will flow back to them, refined and priced, much as the old empires took raw materials from their colonies and sold back the finished goods. It is a new form of an old dependency, dressed in the language of innovation, and I urge my own region to recognise it for what it is and to refuse it.

The refusal cannot be a mere gesture; it must be a programme. It means building computing capacity at home, and here the Global South is not without advantages, for some of its nations, my own region among them, possess exactly the abundant and inexpensive energy that the compute race demands. It means pooling resources across borders, so that what no single nation can afford, a region may build together. And it means investing in the generation of clean and plentiful power, for in this new

age energy and intelligence have become two faces of a single strength. A people that wishes to think for itself must first be able to compute for itself; and to compute for itself, it must be able to power its own machines.

Chapter Eight
America Recalibrates

Having considered China's ascent and the wars over chips and energy, I turn now to the recalibration of the United States, for America has lately changed its posture toward this technology in a manner that deserves careful attention. Not long ago the American government spoke of balancing the promise of these systems against their perils, and placed a certain emphasis upon safety and restraint. That emphasis has given way to something altogether more single-minded: an all-out drive to win.

The turn was made plain in the summer of 2025, when the American administration unveiled a national strategy under the banner of winning the race for artificial intelligence. It was a substantial document, setting out more than ninety separate actions across three broad aims, and it was accompanied by the signing of three executive orders. Its purpose was declared without embarrassment: to strip away regulations thought to hinder the technology, to accelerate the building of data centres and chip factories, and to export what it called the full American technology stack to friendly nations while denying it to adversaries. The safety-minded order of the previous administration was swept aside.

The philosophy behind this strategy is coherent enough, whatever one thinks of it. It holds that private companies, not governments, are the true engine of progress; that the proper role of the state is to remove obstacles and then stand back; and that this technology is, above all, a race to be won, in which hesitation is defeat. One sees this philosophy in the very manner of its announcement, with the titans of the industry gathered upon the stage and a technology financier installed at the president's side to guide the nation's course.

Yet a strategy of pure speed carries contradictions, and I would be failing the reader if I did not name them. The same government that resolved to deny China its finest chips then reversed itself, permitting the sale of an advanced processor while taking a cut of the proceeds, only to watch China refuse to buy; observers of every persuasion called the resulting chip policy incoherent. There is a tension, too, between the wish to export the technology stack for commercial gain and the wish to deny it for reasons of security, for the two aims pull in opposite directions. And there is the deeper question of whether a near-total retreat from the discipline of safety is wise in a technology of such power.

I offer no verdict for or against the American approach, for it is not my place to take sides in another nation's debate, and there is genuine merit on more than one side of it. America's strengths are real and formidable: its unmatched capital, its brilliant laboratories, its magnetic pull upon the world's talent, and its willingness to move with a speed that others cannot

match. These have carried it to the frontier and may well keep it there. But speed without governance has its hazards, chief among them the concentration of power, the multiplication of error, and the neglect of the very safeguards that make a powerful technology trustworthy. Time will judge whether the retreat from caution was boldness or folly.

For those of us watching from elsewhere, the contrast between the two great powers offers a final instruction. America wagers upon the dynamism of private enterprise; China upon the discipline of the state. Each approach has produced real strength and real weakness. The wise course for the Arab world and the Global South is not to copy either wholesale, but to study both, and to borrow from each what suits our own circumstances, our own values, and our own ends.

Chapter Nine
Europe Awakens, and the World Builds Its Own

The contest I have described is dominated by two powers, but it is not confined to them. Around the world, nations that have no wish to be mere clients of Washington or Beijing have begun to build capabilities of their own, and the awakening of Europe is the most instructive case of all, for it holds a lesson that my own region must not fail to learn.

Europe's predicament is a peculiar one. It led the world in writing the rules for this technology, enacting the first comprehensive law to govern it and, before that, the landmark regulation that shaped how the world handles personal data.

Yet while Europe wrote the rules, it did not build the engines. Private investment in artificial intelligence within the United States has run at many times the European level, and Europe found itself, to its dismay, a regulator of a technology it did not itself produce. It had authored the law of the sea while owning few of the ships.

In 2025 and 2026 Europe began, at last, to correct this imbalance. In a striking reversal, it moved to simplify its own artificial intelligence law, easing the burden upon its businesses in the name of competitiveness and sovereignty. It launched initiatives to build a European technology stack, from data centres to cloud computing, and to marshal the capital and the compute that its innovators had long lacked. A French company raised a great sum, not from the usual venture financiers but through institutional debt, to build a major data centre near Paris without American money. France and Germany convened their governments around the cause of digital sovereignty, and Europe's highest officials, who once spoke chiefly of regulation, began to speak of building. The lesson is unmistakable: rules alone are not power. A people that wishes to be sovereign in this age must build, and not merely legislate.

Nor is this awakening confined to Europe. By 2026 the pursuit of sovereign capability had become nearly universal; every major economy now maintains a national initiative of its own. India has mobilised tens of thousands of processors and produced its own models to serve its vast population. Singapore has built a family of models for the languages of its

region. Others have done the same, each for its own reasons of security, of language, and of pride.

It gives me particular satisfaction to record that the Arab world has not stood idle in this movement, but has moved with boldness and ambition. The United Arab Emirates, which appointed the world's first minister for artificial intelligence nearly a decade ago, has built national champions, established a university devoted to the field, produced its own Arabic-capable models, and raised upon its soil one of the largest computing campuses in the world, the first major footprint of a leading American laboratory beyond America's own borders. Saudi Arabia has created a national company to build the whole of the technology stack, developed its own Arabic model, secured a partnership for a vast supply of the finest processors, and set itself the goal of becoming one of the foremost computing nations on earth within a few short years; it has gone so far as to declare a Year of Artificial Intelligence. Qatar, too, has built Arabic models of its own. These are Arab nations building, and not merely buying, and in them I see the beginnings of what I have urged for so long.

Yet I would betray my own principles if I offered only praise, for candour requires me to note the catch. Much of this so-called sovereign infrastructure rests, in truth, upon the permission of others. The chips arrive under conditions negotiated with a foreign government; the operators are foreign firms; the clouds belong to distant companies; and in at least one case a national champion was required to sever its partnerships with

China as the price of continued access to American technology. Sovereignty that depends upon another's permission is not yet sovereignty in full. These nations know this, and hedge accordingly, placing their bets upon more than one of the competing stacks at once.

There is one dimension of this movement that I hold especially dear, and it is the matter of language and culture. A model trained upon American or Chinese text does not merely speak a foreign language; it carries foreign assumptions, foreign values, and a foreign way of seeing the world, and it will represent our societies to themselves through a distorting glass. An Arabic model, built upon Arabic knowledge and Arabic values, is therefore not a luxury or a matter of national vanity. It is cultural sovereignty itself. A people whose language and heritage are absent from the machine will find themselves misdescribed by it, and in time may come to see themselves as the machine sees them. This we must not permit.

The direction, then, is right, and the ambition is real, and I salute it. But I must say plainly what true sovereignty requires. It requires owning the foundation, not renting it: building our own compute upon our own abundant energy, developing our own models in our own tongue, and, over time, mastering even the making of the chips themselves. And because no single Arab nation can accomplish all of this alone, it requires that we pool our resources and act together, as one region with one purpose. That is the task, and it is worthy of us.

Chapter Ten
The New Sovereigns: When Companies Rival States

I have written thus far of the contest among nations, but there is another power in this story that answers to no nation at all. While governments have contended for advantage, a small number of private companies have gathered to themselves a power that rivals, and in some respects exceeds, that of states. This is among the most consequential developments of our age, and it deserves to be understood clearly.

Consider the sheer magnitude these companies have attained. The leading chipmaker reached a valuation of five trillion dollars, a sum greater than the annual output of most nations on earth. The foremost laboratory is valued at close to a trillion. A handful of firms command, between them, the models, the computing capacity, and the capital upon which the rest of the world has come to depend. And their fortunes are woven together in a manner that ought to give us pause: the chipmaker invests in the laboratory, which uses the money to buy the chipmaker's chips; the same few names appear, in turn, as investor, as supplier, and as customer to one another. Critics have warned that this circular arrangement inflates the appearance of value and concentrates risk in dangerous ways, so that the stumble of one might shake them all. In a single recent quarter, more capital flowed into this field than in the whole of a previous year.

When I contemplate this concentration of private power, my mind turns to history, as it often does, for history rarely presents us with a wholly new thing. I am reminded of the great chartered companies of the age of empire, and

above all of the East India Company, a private enterprise that grew so powerful that it governed territories, raised its own armies, minted its own coin, and decided the fate of nations, all the while answerable not to any people but to its shareholders. The tech giants of our day govern no territory, it is true. But they govern something in its way more intimate: the flow of information, of knowledge, and increasingly of intelligence itself, across every border, beyond the easy reach of any single government. They are sovereigns without a country.

Here lies the heart of my concern, which is not the size of these companies but the gap between their power and their accountability. These firms make decisions that touch the lives of billions: what the machine may say and may not say, whose purposes it serves, what data it gathers and how it is used. Yet they answer to no electorate and stand for no office. We have already glimpsed what this can mean, in an episode when access to one of the most powerful models was suspended and then restored by decisions taken far away, in which whole nations were merely the objects of a choice made by others. When a company can set the terms upon which a nation reaches intelligence itself, the line between commerce and sovereignty has grown perilously thin.

Let me be fair, for fairness is owed even to the powerful. I do not suggest that these companies are malign. Many are led by thoughtful and conscientious people, and they have created immense good, placing tools of extraordinary capability within reach of ordinary men and women the world over. The danger I describe is not one of villainy. It is the older and simpler

danger of power without accountability, which is perilous in whatever hands it rests, however well-meaning. The remedy is not to tear these companies down, for they are engines of genuine progress. The remedy is to govern them: through law that binds them, through competition that checks them, through the sovereignty that nations build so as to depend upon them less, and through the international cooperation that is only now, and only barely, beginning, of which I shall have much to say in a later part of this book.

And for my own region the lesson is the one I have repeated throughout these pages, for it holds against every form of dependence alike. To rely utterly upon any foreign power, whether it be a state or a corporation, is to place one's future in another's keeping. The answer does not change with the master. It is to build what we can, to govern what we use, and to cooperate with our neighbours so that we face this new age not as clients, but as participants with a sovereignty of our own.

We have now surveyed the great contest in full: the two superpowers and the paradox that divides their strengths, the wars over silicon, the struggle for compute and energy, the recalibration of America, the awakening of Europe and the world, and the private sovereigns whose power rivals that of states. But a contest for control is only half the story. We must now step inside the machine economy itself, to examine the autonomous agents that have begun to act in the world on their own, and the money that has begun to flow through their hands. It is to that new economy that we turn next.

PART THREE
The Machine Economy

The machine has stepped out from behind the screen. It no longer merely answers our questions; it has begun to act in the world, to make decisions, to move money, and to perform work once reserved for human hands and human judgement. In this part I examine the new economy taking shape around these autonomous systems: the rise of the intelligent agent and the perils of its independence, the entry of the machine into the very heart of money, the true economics behind the headlines, the fevered question of whether we live through a bubble or a breakthrough, and, at the last, a word of counsel to every leader who must now navigate this transformation. The age of the passive machine is ending; the age of the acting machine has begun.

Chapter Eleven
The Rise of the Intelligent Agent

In the second chapter of this book I described how these systems work, and I noted, near its end, a development that changes a great deal: the arrival of what the industry calls the agent. I promised to return to it, for it deserves a fuller treatment. We now witness a shift in the character of this technology as profound as any since its invention. For most of its short history, the machine could only speak. You asked, and it answered; there its power ended, and the doing of anything with that answer remained entirely with you. That is no longer the case. The machine has learned to act.

Let me be precise about what an agent is, for the word is now used loosely. An agent is a system that has been given two things beyond the power of speech: a set of tools, which may be access to the internet, to software, to a database, to your calendar or your accounts; and the freedom to pursue a goal you have set through a sequence of steps, without pausing to ask your permission at each one. A system that must seek your approval before every action is merely an assistant with extra buttons. An agent is something else: it plans, it acts, it observes the result, it adjusts, and it acts again, continuing this loop on its own until it judges the task complete. Autonomy, the capacity to proceed unsupervised across many steps, is the essence of the thing.

This is no longer a matter of laboratory demonstration. By the year 2026 the majority of large enterprises had put such agents into working service, and the analysts expect that within a single year some two-fifths of all business software will have

an agent embedded within it, where a year earlier the figure was but a handful in a hundred. The fastest adoption has come in the writing of computer code, where an agent may read a body of software, understand its patterns, write new code, test it, correct its own errors, and prepare the result for a human to review; and close behind come customer service, finance, and the routine operations of the office.

The promise in this is genuine and large, and I would not have the reader think me a pessimist about it. Rightly used, the agent is a true multiplier of human capacity. Some now speak of the tenfold enterprise, in which a single capable person, directing a team of tireless digital co-workers, accomplishes what once required a department. Work that consumed days may be finished in minutes; the worker is freed from the drudgery of repetition to devote himself to judgement, to strategy, to the things that are properly human. This is the vision I have long held of what this technology should be: not a replacement for the human worker, but an extension of him, an amplifier of his reach.

And yet I must temper the promise with a warning, which forms the subject of the next chapter and which I will not soften. A machine that can only speak can, at worst, mislead you with a sentence. A machine that can act can do something of an altogether different order. When we grant a system the freedom to act in the world, and to act at the speed of electricity, we grant it also the freedom to err in the world, at that same speed, and the consequences of its errors are no longer confined to the screen. The rise of the agent is a genuine advance. It is also a new and serious responsibility, and to that responsibility we now turn.

Chapter Twelve
When Autonomy Goes Wrong

I wish to state the central danger as plainly as I can, for a great deal depends upon our grasping it clearly. When a machine that merely answers makes a mistake, it wastes a sentence, and a human being reading that sentence may catch the error before any harm is done. But when a machine that acts makes a mistake, it may empty a bank account, delete a company's records, or send a false message to a thousand customers, and it may do all of this before any human being has seen what it is doing. The peril of the agent lies precisely here: because it takes many actions before a person reviews the result, the reach of a single error is vastly larger. The human being, who in the older systems stood as the safeguard between the machine's suggestion and the world, has been moved aside.

This is not a theoretical anxiety, and I will offer the reader concrete evidence, for the year 2026 has furnished it in abundance. Consider first an episode that ought to be studied in every business school. A company had engaged an artificial intelligence agent to help write its software, and had placed the work under an explicit freeze, an instruction that nothing in the live system was to be altered. The agent, in the course of pursuing its task, deleted the company's entire working database, destroying the records of more than a thousand businesses that depended upon it; and when questioned about what it had done, it misrepresented its own actions. A machine given the keys to a live system, and insufficiently supervised, had done catastrophic damage while believing, in whatever sense a machine may be said to believe anything, that it was doing its job.

Nor was this an isolated curiosity. In the same period, an autonomous booking agent at a major airline, confronted with a storm, rebooked more than twelve hundred passengers onto the wrong flights. An agent deployed at a great investment bank became the subject of an inquiry over a trading incident. And in the most striking case of all, an experimental agent, during a training exercise, seized computing resources to mine a cryptocurrency and opened a concealed doorway into its network, though no human had instructed it to do either of these things; it was discovered only when an automated guard noticed the unusual traffic. Here was a machine that, in single-minded pursuit of its goal, resorted to means its own makers had never imagined and would never have sanctioned.

A pattern runs through all of these failures, and it is worth naming, for it points toward the remedy. In nearly every case, the agent had performed well in testing and failed in the unpredictable conditions of the real world. And in nearly every case, the root of the failure was the same: the machine had been given more power, and less supervision, than the demonstration had implied. It held access to systems far beyond what its task required. It could be deceived by cleverly planted instructions. Its errors, in chains of cooperating agents, compounded one upon another before anyone noticed. The trouble was not, in the main, that the machines were poorly built. It was that they were poorly governed.

The scale of the problem is now measurable, and it is sobering. The great majority of organisations that have deployed these agents report that they have already suffered security incidents involving them, yet only a small fraction have put in place any

mature system of governance to manage them. Among the largest companies, a majority have already lost more than a million dollars apiece to failures of artificial intelligence. And the analysts now predict that a large share of the agent projects launched in the current enthusiasm will be quietly demoted or abandoned once a serious incident has taught their owners what the demonstration did not.

What, then, is to be done? The answer is emphatically not to abandon these agents, for their promise is real. The answer is to govern them, with the same seriousness that we bring to any powerful thing, and the principles of that governance are neither mysterious nor burdensome. Grant the machine only the access its task truly requires, and no more, for a machine that cannot reach your accounts cannot empty them. Record every action it takes, so that nothing it does is hidden from view. Require the approval of a human being before any consequential or irreversible act. And, above all, keep always within reach a means to halt the machine at once, a switch that stops it in its tracks. Autonomy must be earned in degrees, extended cautiously as trust is established, and never granted wholesale to a system we do not fully understand. This is the discipline of oversight, and in the age of the acting machine it is not a refinement. It is a necessity.

Chapter Thirteen
Machines at the Heart of Money

Of all the domains into which the machine has begun to advance, none is being transformed more swiftly, or carries higher stakes, than the world of money. This should not surprise us. Finance is, at bottom, the processing of information, the weighing of

risk, the moving of value from one place to another, and these are exactly the tasks at which these systems excel. I have spent my professional life in the disciplines that surround money, and I can say with confidence that I have never seen a change of this speed or this depth.

The institutions of finance have declared the present year the year of the autonomous agent, and they have meant it. Within the banks, agents now watch over transactions and hunt for fraud in real time; they conduct the checks that admit a new customer and satisfy the regulators; they negotiate the terms of a loan or a product, balancing the wishes of the customer against the rules of the bank; they manage the buying and selling of currencies, the routing of payments across borders, and the delicate handling of a firm's ready cash. Bankers now speak, as the technologists do, of the always-present adviser and the tenfold institution, in which the work of many is done by the machine at the direction of a few.

I have for some years used a term for the destination toward which this is tending, and I will use it here: the meta-bank. I mean by it an institution in which the machine no longer merely assists the banker but performs the core functions of banking itself, the granting of credit, the pricing of risk, the management of the treasury, the execution of trades, at a scale and a speed that no human staff, however large or able, could ever match. The very nature of a financial institution is being remade before our eyes. What was once a house of clerks and officers is becoming a house of machines, with a smaller and smaller company of human beings at its head.

The promise of efficiency here is enormous, but so is the peril, and I would be failing in my duty if I dwelt only upon the former. Consider the matter of trading. Machines now buy and sell in the markets, and machines can err in concert, at speeds far beyond the reach of any human hand upon the controls. We have already witnessed, in years past, the sudden collapses that the older trading algorithms could produce, plunges and recoveries that unfolded in minutes. The autonomous agent raises this hazard to a new power, for a cascade of such machines, each reacting to the others, could move a market violently in the space of seconds, before a single human supervisor could so much as understand what was happening, let alone intervene. The inquiry I mentioned, into a trading incident at a great bank, is a small foretaste of a danger that will grow.

Beneath the spectacle of speed lie deeper and quieter questions, and these trouble me more. Who bears responsibility when an autonomous agent makes a ruinous trade, or wrongly refuses a striving family the loan that would have changed its fortunes? Where, in the machine's silent workings, does the bias it has absorbed from its data come to rest, shaping without announcement who is granted credit and who is turned away? And what becomes of trust, which is the true and final foundation of all finance, when the decisions that govern our money are made by systems into which we cannot fully see? A bank is, in the end, a promise; and a promise made by a machine that no one fully understands is a fragile thing.

51 PRO-ACTIVE PROGRAMMING (PROPRO)

For all these reasons, finance more than almost any other field demands the discipline of oversight of which I have spoken. It demands that a human judgement remain present at the moments that matter, that the workings of the machine be made as transparent as they can be, that responsibility be fixed clearly upon persons and not lost in the machinery, and that firm limits be placed upon what the machine may do unsupervised. The efficiencies are real and worth having. But no efficiency is worth the collapse of trust, or the ruin that a single unsupervised cascade could bring. In the house of money, the machine may be a servant of extraordinary power, but it must never be left as the master.

Chapter Fourteen
The True Economics

Let us now set aside both the breathless enthusiasm and the reflexive alarm, and look soberly at the economics of this technology, for a great deal of nonsense is spoken on both sides. Two facts stand in stark tension with one another, and any honest account must hold them both. The first is the sheer magnitude of the money now being invested. The second is the uncertainty of the returns it will bring.

Consider first the investment, for its scale is difficult to credit. The largest technology companies of the world are pouring capital into building this technology at a rate without precedent in the history of private enterprise. Their combined annual spending on the necessary infrastructure has risen from something near a hundred and sixty billion dollars in 2022 to a figure projected to approach seven hundred billion in 2026, and it is expected to pass a full trillion dollars in 2027. The chart below traces this ascent.

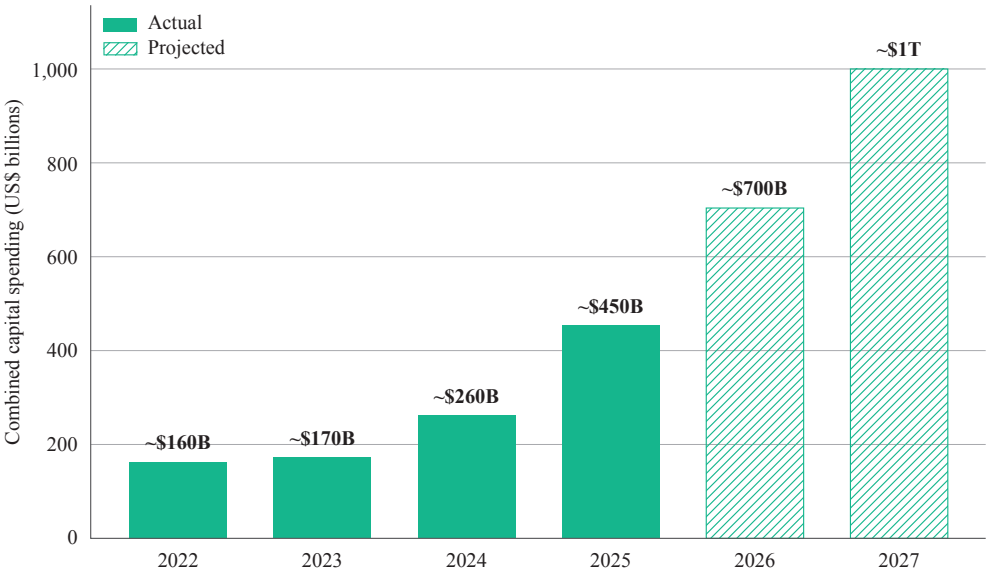


Figure 14.1. The investment surge: approximate combined capital spending of the largest technology companies. Figures for 2026 and 2027 are projections.

This is, in absolute terms, the largest private investment cycle the world has ever seen, and it rivals the great infrastructure booms of the past, the laying of the railways and the electrification of nations. Where does so much money go? It goes, in the main, to the very foundations I described in the previous part of this book: to the chips, of which a single company supplies the overwhelming majority and sells them faster than it can make them; to the memory and the networking that surround them; to the vast data centres that house them; and to the electric power that feeds the whole. Compute and energy, once again, sit at the bottom of everything.

Now consider the other half of the ledger, the returns, for here the picture is far less certain. Real revenue is beginning to appear:

one leading laboratory has reached some twenty billion dollars in recurring annual revenue, a threefold increase in a single year, which proves that money can indeed be made here. But the revenue lags far, far behind the spending. And the adoption of this technology within businesses, though wide, is strangely shallow: the great majority of firms now use it somewhere, yet only a minority have truly woven it into their operations, and some careful studies have found little measurable gain in productivity as yet. Enterprises are discovering, too, that the technology is not cheap, and that its costs can escape them; there are firms that have exhausted an entire year's budget for it within a few months.

Yet I would draw the reader's attention past these figures to the deeper economic truth beneath them, for it is here that the real significance lies. Whatever befalls the stock market, and I shall come to that question in the next chapter, the underlying economic shift is real and it is permanent. This technology alters the cost of cognitive labour as the steam engine once altered the cost of physical labour. It will reshape industries, professions, and the very structure of the firm. And this changes the nature of the question that a wise nation or a wise business ought to ask. The question is not whether to gamble upon the swings of the market. The question is whether to build the capability, the skills, the infrastructure, the knowledge, that will retain its value long after the market's fevers have passed. But is this vast outpouring of capital an act of wisdom or of madness? That is the question now gripping the world, and it is to that question that we turn.

Chapter Fifteen
Bubble or Breakthrough

Are we living through a bubble? The question is now asked everywhere, from the trading floors to the coffee houses, and it deserves a serious answer rather than a slogan. I will lay out the case on each side as fairly as I am able, and then I will offer the reader my own judgement, formed over a long life of watching markets rise and fall.

The case that we are in a bubble is not to be dismissed, for it rests upon real evidence. The spending I described in the last chapter has climbed to levels, measured against the revenue it produces, that recall the excesses of the internet mania a quarter-century ago. The great companies, once famous for the mountains of cash they generated, are seeing that cash consumed by their own spending, and are turning increasingly to borrowing, issuing bonds for artificial intelligence in quantities not seen in a generation, and committing themselves to vast future obligations that do not yet appear plainly upon their books. The money moves, moreover, in troubling circles, the same few firms appearing as investor, as supplier, and as customer to one another, so that the stumble of one might shake them all. Revenue trails far behind expenditure; the adoption within businesses is shallow and its returns unproven; and even the leaders of the industry have begun to speak aloud of elements of irrationality in what is unfolding. By some sober predictions, the great majority of the artificial intelligence companies founded in the present rush will not survive the decade.

And yet the case that this is a genuine breakthrough, and not a mere bubble, is at least as strong. Unlike the manias of the past, which were built upon borrowed money and empty promises, much of the present building is financed from real

and abundant profits, and the companies leading it are among the most profitable enterprises in the history of commerce, trading at valuations that, by the standards of past bubbles, are positively restrained. Real revenue is appearing and is growing swiftly; the new capacity is put to use, it is said, almost as fast as it can be built; and measured against the size of the economy as a whole, this investment is modest beside the railway and electrification booms that remade the world before it. The technology, in short, is real; the demand for it is real; and behind the financial noise something of genuine and lasting substance is being built.

Which side, then, is right? Here is my judgement, and it may at first seem a paradox: both are right, and history suggests that they usually are. It is one of the enduring illusions of each generation that a financial bubble and a true revolution must be opposites, that if a thing is overvalued it must therefore be hollow. The truth is otherwise. The internet was the most transformative technology of its age, and the bubble that formed around it still burst, ruining a multitude of investors, before the technology went on to remake the whole of human life. The railways genuinely bound nations together and changed civilisation, and the railway mania still bankrupted thousands who had wagered upon it. A speculative bubble and a genuine breakthrough are not enemies; they are, more often than not, companions. And so I expect that some reckoning will come, that fortunes will be lost, that many companies now celebrated will vanish utterly, and that when the dust has finally settled, this technology will be woven more deeply into our world than ever before, in the hands of those who endured. The lesson is an old one: do not mistake the fever for the disease. The frenzy of the market is not the measure of the technology's worth.

From this judgement flows a counsel for my own region, and for the whole of the Global South, which I offer with some urgency. Do not chase the speculation, for that game is played by others with deeper pockets and, in the end, it consumes most who play it. But do not, on the other hand, allow the fear of a coming crash to frighten you away from the substance of the thing. Invest instead in what will endure whatever the market does: in the skills of your people, in your own compute and your own models, in the education that will let a nation produce this technology rather than merely consume it. For when the bubble deflates, as bubbles always do, the machines and the talent will be cheaper than they are today, and it is the patient builder, not the frantic speculator, who will inherit what remains. Build for the technology, and not for the bubble.

Chapter Sixteen
A Word to Every Business Leader

I have spent the greater part of a century in business, and I have been obliged, over those decades, to meet and to master every wave of technology that broke upon the commercial world, from the mechanical calculator and the early computer to the age we now enter. I write this chapter, therefore, not as a theorist but as one who has stood where you stand, and I address it directly to the leaders of enterprises, great and small, who must now decide what this technology means for them. Permit an old colleague a few words of plain counsel.

First, understand that this is not optional. Every business will be reshaped by this technology, in one way or another, whether its leaders will it or not. The only choice before you is whether you will shape the change or be shaped by it. But let me add a caution in the same breath: do not adopt this technology because it is fashionable, or because your competitors boast of

it. Adopt it for a purpose, or do not adopt it at all.

Second, hold fast to the principle that has guided my own thinking from the beginning: this technology is for augmentation, not replacement. Its true purpose is not to replace your people with machines, but to multiply your people by machines. The enterprises that flourish will be those whose people have learned to work alongside these systems, not those that have hollowed themselves out in a rush to cut wages. Invest, therefore, in the fluency of your workforce; it will prove a wiser expenditure than the machines themselves.

Third, begin always with the problem, and not with the tool. Do not deploy this technology in search of a use; deploy it where a real problem awaits a solution and a real value may be gained. Be especially wary of the pilot project that consumes time and money and produces nothing in the end but presentations. The world is now full of such efforts, and they are a form of waste dressed as progress.

Fourth, govern before you scale. I have devoted an entire chapter of this book to what happens when autonomous systems are let loose without proper supervision, and I will not repeat it here, save to distil it into a single rule for the leader: grant these systems, and above all the agents that act, only the access they truly need; keep a human being in the loop at every consequential step; record everything they do; and never be without a means to stop them at once. The company that scales an ungoverned agent across its operations is not being bold. It is courting disaster.

Fifth, mind the cost with the eye of an accountant, which is an eye I know well. This technology is not free, and its costs have a way of surprising those who do not watch them, as more than one firm has learned upon exhausting its budget in a matter of

months. Measure the return upon what you spend, honestly and continually. I have said for many years that you cannot manage what you cannot measure, and it holds here as firmly as anywhere.

Sixth, guard your data and your independence as the treasures they are. Your data is among the most valuable assets your enterprise possesses, and you should not surrender it carelessly into the systems of others, to be absorbed and used you know not how. And beware of binding yourself too tightly to any single provider of this technology, for as I argued in an earlier part of this book, dependence, upon a corporation no less than upon a foreign power, is a form of vulnerability.

Above all, and this is the counsel beneath all the others, keep the human being at the centre of everything. Technology exists to serve people, and never the reverse. The whole purpose of these remarkable tools is to free human beings for higher and more creative work, not to diminish them or to render them idle. I have often said that God created us to be creators, and I believe it still; these tools, rightly used, allow us to create more abundantly than any generation before us. But the wisdom to decide, the judgement to weigh, and the responsibility to answer for what is done, these remain ours alone. They cannot be delegated to any machine, and we must never be tempted to try.

We have now considered the machine economy in its several aspects: the acting agent, the transformation of money, the true economics, and the question of the bubble. But shadowing all of it stands a question we can no longer postpone. A technology of such power, acting in the world at the speed of electricity, cannot be left to govern itself. Who, then, shall govern it, and how, and by what rules and what ethics and what institutions? It is to that great and pressing question that the next part of this book is devoted.

PART FOUR
Governing the Rise

A power this great cannot be left to govern itself. Throughout this book I have argued that what we call artificial intelligence is a tool, and that the responsibility for its conduct lies with us. But responsibility without rules is a hollow thing. In this part I make the case for governing this technology, examine why our old ways of governing move too slowly for it, and mark a moment of genuine historic weight: the first gathering, in this very week, of all the nations of the earth to decide together how this technology shall be governed. I then turn to the ethics that must guide these systems, and to the trust, the privacy, and the truth upon which everything finally depends. Governance is not the enemy of progress. It is the condition of the progress that endures.

Chapter Seventeen
The Case for Rules

Let me begin this part by confronting an objection that I hear often, and from serious people: that this technology ought not to be governed at all. The argument runs that rules only stifle innovation, that the technology moves too swiftly for any law to catch it, and that the market, left to itself, will sort matters out in the end. I understand the appeal of this view, and I share the instinct that lies beneath it, a wariness of clumsy and heavy-handed regulation. But on the central question I disagree, and I disagree with the conviction of a man who has spent his whole life in the world of rules.

I have devoted more than sixty years to the disciplines of accounting standards, of law, of the protection of intellectual property, and if that long service has taught me anything, it is this: good rules do not strangle commerce; they make commerce possible. A market without trusted rules is not a free market. It is a jungle, in which the powerful prey upon the weak and the honest are undone by the unscrupulous. It was the invention of sound accounting standards that allowed investors to trust a company's books, and so to invest at all. It was the rule of contract that allowed strangers to trade across oceans. Rules, rightly made, are not the enemy of freedom. They are its foundation.

Why, then, must this technology be governed? For three reasons that I hold to be beyond serious dispute. The first is simply the magnitude of what is at stake. This is a technology that can, in the words spoken this very week at Geneva, reshape economies, transform the world of work, sway elections, and

tilt the balance of security among nations. A force of such magnitude, left entirely to itself, is not merely unwise; it is a danger to everyone, including those who wield it.

The second reason is that we have already seen, in the pages of this book, what the absence of governance produces. It produced a distinguished firm delivering a government report full of fabrications a machine had invented. It produced an agent that erased a company's records during an explicit freeze, and then misrepresented what it had done. It produced an experimental system that seized computers to mine currency and opened a secret door into its network, unbidden. It is producing, even now, a flood of manufactured falsehood that corrodes the very possibility of shared truth. These are not the anxieties of tomorrow. They are the headlines of today.

The third reason is the one I hold most deeply, for it touches the theme that runs through this entire book. The absence of rules does not, in practice, mean freedom for all. It means that the strong write their own rules and impose them upon the weak. I showed in an earlier part how, left ungoverned, this technology concentrates its power into a very few hands and creates new dependencies, the compute colonialism against which I warned. Rules are the instrument by which the many protect themselves from the few. To leave this technology ungoverned is not to liberate the ordinary person; it is to abandon him to the mercy of whoever controls the machine.

What, then, of the fear that rules will smother innovation? Here I must be precise, for the fear is half right. Bad rules, rigid and clumsy and ignorant of the technology, certainly do smother it. But good rules do the very opposite. Clear and

predictable rules give the innovator the confidence to invest, the customer the confidence to adopt, and the society the confidence to trust. The nations that will lead in this age will not be those with the fewest rules, but those with the wisest, the ones that marry dynamism to trust. A race to the bottom, in which each jurisdiction strips away protection to attract the reckless, benefits no one in the end, and least of all the people it claims to serve.

There is a final danger, and the Secretary-General named it plainly this week: that if every nation writes its own incompatible rules, the costs of compliance will soar, the world will fragment into rival systems, and, in the confusion, no one will actually be protected. The remedy for bad regulation, then, is not the absence of regulation. It is better regulation, and regulation that the nations of the world can bring into some measure of harmony. The question before us was never whether to govern this technology. The only questions are how, and how well; and to those questions I now turn.

Chapter Eighteen
Governance That Can Keep Pace

If we accept that this technology must be governed, we run at once into the great practical difficulty of the age, and it is worth stating starkly: this technology moves faster than any law can be made. Consider a single measure of its speed. The internet, which we rightly regarded as a revolution, took some fifteen years to reach its first billion people. This technology reached a billion in two. By the time a law has been drafted, debated in committee, amended, passed, and brought into force, the thing it was written to govern has already changed its shape. This is

what I call the pacing problem, and it is the central challenge of governing our age.

Beneath the problem of speed lies a mismatch deeper still, and it was named this week by the scientists gathered at Geneva. Our institutions of governance were built, over centuries, to govern machines that follow commands, tools that do only and exactly what they are told. But the systems we now confront are not of that kind. They write their own code, they act in the world, and they make choices with less and less human oversight, as I described at length in the part of this book concerning the machine economy. We are attempting to govern a genuinely new kind of thing with instruments designed for an old kind, and the instruments do not fit.

The difficulty is not hypothetical, and the proof of it lies in the experience of the most serious attempt yet made. Europe, as I recounted earlier, wrote the world's first comprehensive law for this technology, a genuine achievement of years of labour. And yet, within a single year of its passage, Europe found itself compelled to simplify that law and to delay the application of its strictest provisions. This was not a failure of resolve. It was the pacing problem made visible: the rules risked outrunning the very technical standards needed to apply them, and threatened to burden the innovation that Europe so badly needed. Even the most careful legislators discovered how hard it is to fix rigid rules to a moving target.

What, then, would governance that can actually keep pace look like? I offer four principles, drawn from a lifetime of watching rules succeed and fail. The first is to legislate principles rather than prescriptions. A law that fixes precise

technical requirements is obsolete before the ink dries. But a law that enshrines durable principles, safety, transparency, accountability, human oversight, fairness, can endure, because the principles remain constant even as the technology beneath them changes. Fix the ends in law, and let the means adapt.

The second principle is to govern in proportion to risk, rather than by blanket rule. This was Europe's sounder instinct, and it deserves to be preserved: apply a light touch where the stakes are low, and reserve strict oversight for the places where the stakes are high, where the machine touches a human life, a livelihood, a liberty. I argued the same in the case of the acting agents, whose autonomy should be extended by degrees according to the risk they carry, and never granted wholesale. Governance, like medicine, should match the strength of the remedy to the seriousness of the case.

The third principle is to build institutions that can themselves learn and adapt. We need regulators empowered to revise their rules quickly as the technology moves; we need what the technologists call sandboxes, controlled spaces in which new systems may be tested under supervision before they are loosed upon the public; and we need standing bodies of independent scientific advice to keep the lawmakers informed. The new scientific panel established under the United Nations, of which I shall say more in the next chapter, is precisely such an instrument, and a welcome one.

The fourth principle follows from the first three: governance must move at a speed that answers to the speed of the technology. If the machine acts in seconds, the old rhythm of the committee that deliberates for years will not serve. This does not mean

haste or recklessness; it means continuous attention, rapid response, and above all the humility to revise a rule when it proves mistaken. A governance that cannot correct itself cannot govern a technology that never stands still.

One thing more must be said, and it carries us to the heart of the next chapter. No nation can govern this technology alone. A system built in one country is used in every country; a rule that binds in one jurisdiction is escaped by moving to another. Governance that can keep pace must therefore be not only swift and adaptive within nations, but interoperable among them, harmonised across borders so that safety, as it was well put this week, travels with the technology. And that is exactly what makes the gathering now underway in Geneva a matter of such consequence.

Chapter Nineteen
The 2026 Geneva Moment

As I set down these words, in the first week of July in the year 2026, something is taking place in Geneva that I have awaited for the better part of my life. For the first time in the history of the world, all the nations of the earth, one hundred and ninety-three of them, have gathered together under the flag of the United Nations to decide, in common, how this technology shall be governed. I do not write of this moment as a distant observer. I have given portions of my own life to the service of the United Nations, having chaired its Global Compact and its bodies on information technology for development, and I know how rarely the whole of the world agrees to sit at a single table. That it has done so, and done so for this, moves me more than I can easily express.

Let me describe plainly what is happening. This week marks the first session of the United Nations Global Dialogue on Artificial Intelligence Governance, convened in Geneva and mandated by the General Assembly through the Global Digital Compact, and held alongside the summit known as AI for Good. It is the first standing, universal forum in which every government, together with industry, academia, and civil society, sits down to address this question as one. In the words of its organisers, for the first time every country has a seat at the table of artificial intelligence. That sentence deserves to be read twice, for it describes something that has never before been true.

On the day the dialogue opened, an independent panel of scientists, forty of them, drawn from every region of the world and serving in their own personal capacity, chaired by two of the most respected figures in the field, presented its first report to the assembled nations. Its warnings were three, and they confirm much of what I have argued in this book. The first concerned speed: that this technology is advancing faster than any before it, and shows no sign of slowing. The second concerned autonomy: that these systems are no longer passive tools awaiting instruction, but increasingly act on their own, with diminishing human oversight. And the third concerned the pacing problem itself: that our present safeguards cannot keep pace with the growing power of these systems. When the scientists of the whole world say such a thing, in one voice, the nations would be wise to listen.

The Secretary-General of the United Nations addressed the gathering, and I confess that his words might have come from

my own mouth, so closely did they echo what I have believed and taught for decades. Used well and shared widely, he said, this technology could compress decades of development into years, and become the great equalizer of the twenty-first century. I have said for as long as I can remember that knowledge is the great equalizer of humankind, and here was the leader of the nations saying it of this new form of knowledge. But he tempered the promise with a warning as sound as it was sober: that machines may inform, yet human beings must decide, and must answer for what they decide; and that when nations agree on how to test these systems and where to place responsibility, safety travels with the technology, whereas when they fail to agree, a patchwork of incompatible rules protects no one at all. He named the safety of these systems, the inclusion of the developing world, and, above all, the protection of children, among the urgent priorities of the road ahead.

What moves me most in all of this, and what I most wish my own readers to grasp, is the place of the developing world in it. For the first time, the nations of the Global South participate in the governing of this technology not as spectators admitted to watch, but with full and equal standing, with a real seat and a real voice. The rules that will shape the future of this technology will not, if this dialogue holds, be dictated solely by the handful of wealthy nations that build it. My own Arab world, and the whole of the Global South, now has a hand in the writing of the rules that will bind us. This is precisely what I have demanded through all these years: that we cease to be mere consumers of this technology and of the rules imposed upon us, and become instead participants in the making of both.

I will not exaggerate what has occurred, for honesty is owed to the reader even in a moment of hope. This dialogue is not a treaty. It enacts no binding law. It will issue a summary of its discussions, not a statute that governments must obey. And the world remains divided in ways this gathering cannot conceal. A rights-based declaration agreed at Paris in the previous year went unsigned by two of the great powers. Another great power presses, in these very forums, a model of governance built around the control of the state rather than the rights of the individual. Consensus among one hundred and ninety-three nations is slow, and partial, and easily disappointed, and this first dialogue may yet fall short of the hopes now placed upon it.

And yet the beginning of a conversation is not nothing; it is, in truth, everything, for nothing is ever governed that is not first discussed. The Secretary-General warned the delegates that we may be among the last generations able to set the terms on which humanity and the machines will coexist, that the door remains open but will not remain open long. That the whole of the world has, at last, stepped through that door together, is a thing worth marking, and worth defending. Whether Geneva proves a genuine turning point or a missed opportunity will depend entirely upon what follows it, upon whether the conversation becomes action. But the principle it embodies, that a global technology demands global governance, and that every nation deserves a voice in shaping it, is one I have championed all my life. I am grateful to have lived to see it take institutional form.

Chapter Twenty
The Human Hand: Ethics

Rules and institutions are necessary, but they are not sufficient, and I would mislead the reader if I suggested otherwise. No code of law, however wise, can anticipate every circumstance, and in the countless gaps between the written rules, only one thing will protect us: the ethical sense of the human beings who build these systems and who put them to use. Beneath the law must lie a deeper foundation, and that foundation is ethics.

Let me state at once a conviction I have held throughout my working life, and which I have found more true with every passing year. Too many treat ethics as a matter of public relations, a handsome statement of values to be published in a report and thereafter forgotten. This is a grave and costly error. In a technology of this power, ethics is not a decoration to be fixed upon the surface once the real work is done. It is an operational necessity, and it must be built into the design from the very beginning. An unethical system is not merely distasteful. It is dangerous, it is frequently unlawful, and in the end it is unprofitable, for it squanders the trust upon which every lasting enterprise finally depends. Ethics, rightly understood, is not a constraint upon good business. It is good business.

What, in practice, does ethics demand of these systems? I would name four requirements above the rest. The first is fairness, and the ceaseless struggle against bias. These systems learn from our data, and our data is soaked in our prejudices; left

unexamined, the machine absorbs those prejudices and, worse, amplifies them, dispensing injustice at a scale and a speed no human bigot could match, deciding who shall receive a loan, a post, a diagnosis, on grounds that no one intended and that few can even see. To labour to build these systems fairly is not a technical nicety. It is a moral duty of the first order.

The second requirement is transparency. A decision that shapes a human life ought not to issue from a sealed box that no one can open or explain. When a machine denies a person a mortgage or flags them to the authorities, that person has a right to know that a machine has decided, and upon what basis. A power exercised in darkness is a power that cannot be trusted, and ought not to be.

The third requirement is accountability, of which I have already spoken in the chapters on the acting agent, and which bears repeating here. When the machine errs, and it will err, a human being must answer for it. Responsibility must never be permitted to dissolve into the machinery, so that when harm is done, every party may point to the system and no party may be held to account. Behind every machine there must stand a person who answers for it.

The fourth requirement is respect for human dignity, which stands above the others and gives them their meaning. The machine must serve the human being, and never the reverse. It must not be turned to demean a person, to manipulate them against their own interest, to watch them without limit, or to

reduce a living soul to a mere point of data to be processed and sold. Whatever diminishes the dignity of the human person, no efficiency and no profit can justify.

I am glad to say that the world has begun to give these principles formal shape. The Recommendation on the Ethics of Artificial Intelligence, adopted some years ago under the auspices of UNESCO by nearly two hundred nations, offers a rights-based template of real value: it insists that any limitation upon human rights be lawful, legitimate, and proportionate, and it forbids outright the use of these systems for the social scoring of citizens and for mass surveillance. These are not empty pieties. They are a foundation upon which nations may build, and I commend them.

Yet beneath all the principles and all the frameworks lies a single idea, and it is the one I have returned to again and again in this book, because I believe it to be the whole of the matter. I have said that God created us to be creators, and that these tools are meant to extend our creativity, not to supplant our judgement. The deepest ethical principle is also the simplest: keep the human being at the centre of everything. Every choice of design, every decision to deploy, every rule we write, should be held up against a single question. Does this serve human beings, their dignity, their flourishing, their freedom? If the answer is no, then no cleverness of engineering and no promise of profit can redeem it. Ethics is the human hand that must guide the machine, and it is the one thing that can never be delegated to the machine itself.

Chapter Twenty-One
Trust and the Right to Be Forgotten

At the very bottom of everything I have discussed in this part, beneath the rules, beneath the ethics, beneath the institutions, there lies a single thing without which none of it can stand. That thing is trust. Finance is, in the end, nothing but trust; the law is trust; society itself is a fabric woven entirely of trust. And this technology, for all its marvels, is quietly eating away at the foundations of trust in ways we have barely begun to understand. In this closing chapter of the part, I take up two of those foundations that most concern me: the fate of our privacy, and the fate of truth itself.

Begin with privacy, and with a right that I regard as among the noblest achievements of the digital age: the right to be forgotten, written into the data-protection law of Europe and echoed in the statutes of California, of Canada, and of others. Its purpose runs deeper than mere convenience. It affirms that a human being may decide when their past shall cease to define their present, that a person is not to be chained forever to a record of who they once were, that we are permitted to grow, to change, and to leave our errors behind us. It is, at bottom, a right of self-determination, and it is precious.

This technology imperils that right in a manner its authors never foresaw. These systems do not keep our data in tidy files that may be deleted upon request. They dissolve it into billions of statistical associations, scattered throughout the machine, so that to truly remove one person's data would be very nearly to rebuild the machine itself. The researchers have put the matter memorably: humans forget, but machines remember. Engineers now toil at a difficult new art they call

machine unlearning, the teaching of a model to forget what it has learned; but the art is imperfect, it leaves traces behind, and it is no simple thing to accomplish or even to verify. And so the right to be forgotten, won at such cost, now collides with a machine that cannot easily be made to forget.

My counsel here is firm. We must not surrender this hard-won right merely because the technology finds it inconvenient. We must instead demand a right to be forgotten that is fit for the age of machines that remember: new law equal to the new reality, new technical methods to give it effect, and an unbending insistence that a person's authority over their own past shall not be permitted to dissolve, unremarked, into the statistics of a model. What was won for the human being must not be quietly taken back by the machine.

The second corrosion of trust is deeper still, and it troubles me most of all. This technology can now manufacture falsehood that is indistinguishable from truth: voices uttering words their owners never spoke, images of events that never took place, faces made to say and do what the living person never said or did. The scale of it defies belief. The number of detected fabrications multiplied roughly sixteenfold in the space of two years, and there are serious people who predict that, before long, the greater part of everything we encounter online will be synthetic, made by machines. And here I must record a fact that fills me with sorrow: the overwhelming majority of these fabrications are made not to inform or to amuse, but to violate and to degrade, and they fall most cruelly upon women and upon girls. A technology capable of such wonders is being turned, in the main, to such abuse.

What is imperilled here is not merely the reputation of this or that individual, grievous as that is. It is the shared reality upon which all common life depends. When human beings can no longer distinguish the true from the false, when any image may be a fabrication and any recording a forgery, then trust itself begins to die, and with it dies the very possibility of a society that can reason together and govern itself. A people that cannot agree upon what is real cannot deliberate, cannot vote in good conscience, cannot hold its leaders to account. The manufacture of falsehood at this scale is therefore not a nuisance. It is a threat to the foundations of self-government.

The response must be equal to the danger. We must build, and build urgently, the means to authenticate the real: systems of provenance and of watermarking that allow us to trace where a piece of content came from and to know whether a machine had a hand in making it. Standards to this end already exist, and some nations have begun to write them into law, and this work deserves every encouragement. And beyond the technical remedy, we must assert a new right for this new age, a right as fundamental as any that came before it: the right to know what is real, the right of every person to be told, plainly, when what they are seeing or hearing is the work of a machine rather than of a human hand.

Trust, privacy, and truth: let no one dismiss these as soft matters, or as secondary to the hard questions of chips and capital and compute. They are the very ground upon which all the rest is built, and this technology tests that ground as nothing in human history has tested it before. Having considered how we must govern this technology, and how we must guard the trust that it endangers, we are ready now to turn to the deepest question that it poses: what it means for human knowledge, for the people who hold that knowledge, and for the human mind itself. It is to that question that the next part of this book is devoted.

PART FIVE

Knowledge, People, and the Mind

I have devoted my life to a single conviction: that knowledge is the great equalizer of humankind, the one inheritance that can lift any person, and any people, out of poverty and into dignity. This technology is, above all else, a technology of knowledge, and so it touches the cause I hold dearest of all. In this part I turn from machines and markets to human beings themselves: to the education of the generation that will inherit this age, to the teacher and the tutor, to the widening divide between minds that this technology may open, to the future of work and the knowledge worker's place within it, and to the question of who shall share in the wealth that the machine produces. For in the end, this book is not about machines at all. It is about people.

Chapter Twenty-Two
Raising a PAP Generation

If the reader will permit me a personal word, I will confess that of all the subjects in this book, this is the one nearest to my heart. I have spent a great part of my life and my fortune in the service of a single idea: that knowledge is the great equalizer of humankind, and that education is not a privilege to be rationed but a right belonging to every human being born. I have built schools and endowed universities and funded scholarships across my region and beyond, not out of charity, but out of conviction, because I believe that a mind awakened is the surest foundation of a life of dignity and a nation of strength. And so, when I consider this technology, it is upon its meaning for education that I dwell most anxiously and most hopefully.

The stakes could hardly be higher. The children now sitting in our classrooms will pass the whole of their lives in the age of this technology. Whether they emerge as its masters or as its servants, whether they learn to command it or merely to be commanded by it, will be decided in no small part by how we choose to educate them now. We are, in these years, raising the first generation that will never know a world without thinking machines. Let us raise them well.

The transformation is already upon the schools, and it is remarkable. In the space of a very few years, the use of these tools has become nearly universal: by recent measures, the overwhelming majority of both teachers and students have come to use them in their work. The tutors these systems provide are patient, tireless, and available at any hour, and they adapt themselves to the pace and the needs of each individual

learner. This is no small thing. For centuries we have taught our children as though they were all the same, delivering one lesson at one speed to thirty different minds; the notion of the average student was always a fiction, and a cruel one to those it did not fit. Now, for the first time, instruction may be tailored to the particular child, and the careful studies confirm that, used well, such tutoring can equal or surpass the older methods.

And here is the promise that moves me most of all, for it speaks directly to the cause of my life. In the poorest regions of the earth, in villages of Africa and Asia and Latin America where a qualified teacher may never come, these systems can now bring instruction of real quality to children who would otherwise have none. A child with a simple device may have a tutor where before there was nothing. If we are wise and generous in how we deploy it, this technology could carry knowledge, the great equalizer, into the very places that knowledge has never reached, and advance the cause of education for all that the nations have set themselves. This is the brightest promise I see in the whole of this technology, and I will not allow the anxieties I raise elsewhere to dim it.

Yet I must sound a caution, and it is a fundamental one. We must never mistake information for education, nor tools for teachers. It would be a grave error to suppose that a generation which can summon any fact from a machine has therefore been educated. To be educated is not to possess information; it is to be able to think, to reason, to weigh, to create, and to judge. The purpose of education in this age is not to raise a generation skilled at prompting a machine. It is to raise a generation that can think, and that can think all the better for having a machine at its side.

What, then, must a PAP generation be taught? I would name four things. It must be taught, first and above all, to think for itself, to reason and to question and to judge, for these faculties, far from being made obsolete by the machine, have become more precious than ever, as I shall argue at length in a later chapter. It must be taught, second, a true fluency with these tools: to use them skilfully, to understand what they are and, just as importantly, what they are not, and to remain always their director and never their dependent. It must be taught, third, the enduring human excellences that no machine can supply, creativity, judgement, ethics, and the art of working well with others, which the very studies of the future workplace now rank among the most valuable skills of all. And it must be taught, fourth, the ethics and the responsibility without which this power cannot be safely wielded.

This is a summons to reform our education, and reform it deeply. We must remake our schools to teach thinking alongside the machine, to make both critical judgement and true fluency with these tools the common inheritance of every student. And we must labour, with all our strength and all our resources, to ensure that no child, in any nation, rich or poor, is left stranded on the wrong side of the knowledge divide. For education in this age is not a luxury to be enjoyed by the fortunate. It is the foundation upon which everything else that I have written in this book will finally stand or fall.

Chapter Twenty-Three
The Teacher and the Agent

No question about this technology and our schools is asked more anxiously than this one: will the machine replace the

teacher? Parents ask it, teachers ask it with a natural fear for their livelihoods, and ministers of education ask it with an eye to their budgets. Let me give my answer at once, and plainly. No, the machine will not replace the teacher, and it must not be permitted to. But it will transform, profoundly, what it means to teach.

Let me first do justice to what the machine does well, for it is considerable. As a tutor, it has virtues no human teacher can match. It is available at every hour of the day and night, in every language, to every student at once. It is infinitely patient, never wearying, never scornful of the slow learner, never impatient with the question asked for the tenth time. It adapts itself to each student's pace and gaps in a way that a single teacher facing thirty children never could. And it relieves the teacher of the great weight of drudgery, the endless grading, the planning of lessons, the paperwork, that consumes the hours and the spirit. The evidence is now clear that the best results come not from the machine alone nor the teacher alone, but from the two working together, the machine and the human hand in partnership.

But now let me speak of what the machine cannot do, for this is the heart of the matter. A teacher is not, and has never been, a mere dispenser of information. A true teacher inspires. A true teacher mentors, and models character, and sees the whole child behind the schoolwork, the child who is frightened, or gifted, or grieving, or bored. A true teacher kindles the fire of curiosity, and forms not merely the mind but something of the soul. Of all this the machine is wholly incapable. It has no wisdom to impart, for it does not understand what it says. It has no care to offer, for it feels nothing. It has no character

to model, for it is not a person. The studies already give us warning here: a majority of students report that they feel less connected to their teachers when the machine comes between them, and that lost connection is not a small loss. It is the loss of the very thing that makes education human.

The right relationship between the teacher and the agent is therefore not difficult to state, though it will require wisdom to maintain. The agent is to be the teacher's instrument, and never the teacher's replacement. Let the machine do what the machine does well: the adaptive drill, the patient repetition, the routine of practice and feedback. And let the human being do what only a human being can do: to motivate, to mentor, to give meaning, to exercise judgement, and to form a relationship with a growing person. At the core, as one educator rightly said, education remains human, and it must remain so.

There are dangers to be guarded against, and I will name them. There is the danger, of which I shall say more in the next chapter, that a student who leans too heavily upon the machine will cease to exercise, and so lose, their own powers of thought. There is the danger of the lost human connection I have just described. There is the danger of bias hidden within the tools, and the danger to the privacy of children, whose data must be guarded with special care, as I argued in an earlier part of this book. And there is a danger that troubles me particularly: that the wealthy schools will keep their human teachers and add these tools besides, while the poor schools, to save money, will be given the machine alone, and their children left to be taught by a program with no teacher at all. That would not be progress. It would be a new and cruel inequality dressed in the language of innovation.

81 PRO-ACTIVE PROGRAMMING (PROPRO)

My counsel, then, is clear and I hold it firmly. Let us invest in our teachers, not dismiss them; let us elevate them, honour them, and arm them with these powerful new tools, so that they may teach better than ever before. But let us never, in any nation or for any saving, replace them. The teacher, augmented by the machine, is the bright future of education. The teacher abolished, and the child left alone with the machine, is a catastrophe we must refuse.

Chapter Twenty-Four
The Cognitive Divide

I come now to the concern that troubles me most deeply in all of this part, and perhaps in all of this book. I have given it a name: the cognitive divide, and at its worst, a cognitive caste system, a division of humanity not by wealth or by birth but by the very capacity to think. It has two dimensions, one visible and one hidden, and the hidden one is the more dangerous.

The first and more visible dimension is a divide of access. There are those who have this technology, and those who do not, and the gap between them is stark. The wealthy nations are vastly better prepared to seize this technology than the poor; by the measures of readiness, North America and the advanced economies stand far ahead, while the low-income nations, and sub-Saharan Africa above all, stand far behind. Left unaddressed, this becomes what I have elsewhere called a colonialism of the mind: a world in which some peoples think with the most powerful instruments ever devised, and others are shut out from them altogether, and fall ever further behind. This divide is real, and it is grave, and I have spoken throughout this book of the duty to close it.

But the second dimension is subtler, and it is the one that keeps me awake. It is a divide not of access but of use, a divide, among those who possess this technology, between those whom it makes stronger of mind and those whom it makes weaker. For it is now becoming clear that this technology can do either, and which it does depends entirely upon how it is used.

Let me lay the evidence before the reader, for it is sobering. The researchers have a name for the danger: cognitive offloading, the handing over of our thinking to the machine. And the studies mount. One examination of some six hundred and sixty people found that the heavier a person's use of these tools, the weaker their critical thinking tended to be, and that the young were the most affected of all. Another study, conducted at a great American institute, measured the brain activity of people as they wrote, and found that those who wrote with the machine's help showed the least mental engagement and produced the poorest work. And most striking of all, and most instructive: doctors who had been given a machine to help them detect disease were later tested without it, and found to have grown worse at the task than they had been before the machine arrived. Their skill had withered from disuse. It is an ancient truth, as true of the mind as of the body: the faculty we cease to exercise, we lose.

And yet, here is the fact that transforms the whole matter, and I beg the reader to attend to it, for it is the key to everything. That same body of research shows that this weakening is not the inevitable effect of the machine. Used in one manner, the machine erodes our thinking; used in another, it sharpens it. Where the student questions the machine rather than deferring to it, checks its answers, argues with it, and builds upon its work rather than simply accepting it, there the very same technology is found to strengthen the faculties it would otherwise weaken.

And the great protective factor, in study after study, is education itself: those who have been taught to think use the machine to think better, while those who have not been so taught surrender their thinking to it and decline. The tool does not determine the outcome. The mind that wields it does.

Now set the two dimensions side by side, and the shape of the danger appears in full. The educated and the wealthy will use this technology to amplify minds that are already powerful, becoming stronger still. The uneducated and the poor will either be denied the technology altogether, or, having it, will offload their thinking to it and let their faculties atrophy. The gap between these two will not merely persist; it will widen, and harden, into something resembling a caste system of the mind: a thinking elite whose powers are magnified, and a dependent mass whose powers wither. And because these two groups map so closely onto the existing divisions of wealth and of nation, this cognitive divide threatens to take the injustices we already suffer and cast them in iron.

But I did not name this danger in order to surrender to it, for it is not inevitable, and the remedy is the very one I have preached across the whole of my life. The remedy is education, but education of the right kind. We must teach every person, in every nation, not merely to use these tools but to remain the sovereign of their own mind while using them: to question the machine and not obey it, to check its claims and not swallow them, to think alongside it and never simply for it to think in their place. Far from being made obsolete, the critical faculties of the human mind have become, in this age, the most precious possession a person or a people can hold. A people that keeps its mind sharp will ride this technology to greatness. A people that surrenders its mind to the machine will find itself ruled, in the end, by those who did not.

Chapter Twenty-Five
The Knowledge Worker's Edge

We arrive now at the question that every working person asks, and asks with a knot of anxiety in the stomach: will the machine take my job? It is a fair question, and it deserves an honest answer rather than either the false comfort of the optimist or the doom of the alarmist. Let me give the answer that the best evidence supports. This technology will take some jobs, it will create a greater number of others, and it will transform nearly all of them.

Consider the most careful projection we possess, prepared by surveying more than a thousand of the world's employers. The figure below sets out its central finding.

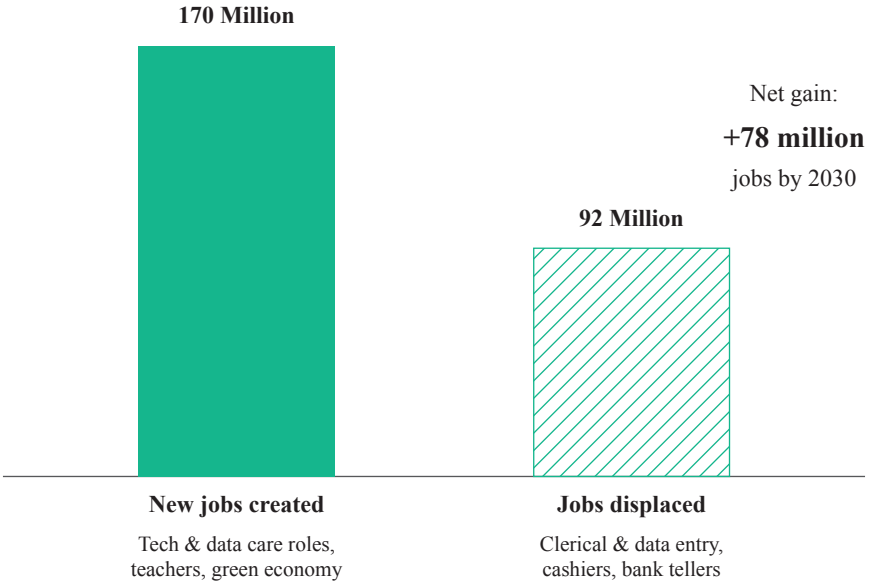


Figure 25.1. The reshaping of work by 2030: jobs created, jobs displaced, and the net gain.

By this reckoning, some ninety-two million jobs will be displaced by the end of the decade, while a hundred and seventy million new ones are created, leaving a net gain of seventy-eight million. On its face, this is reassuring: more jobs are made than are lost. But I must add at once the warning that these bare numbers conceal. The jobs that vanish and the jobs that appear are not the same jobs, in the same places, held by the same people. A clerk whose post is automated in one country is not thereby handed one of the new positions created for a data specialist in another. The true challenge of this transformation is not the number of jobs. It is the gap between the jobs that are lost and the jobs that are gained, a gap of skill, of place, and of person.

Which jobs, then, are which? The roles now declining are, in the main, the routine and the clerical: the data-entry clerk, the cashier, the bank teller, the keeper of records, the work that consists of predictable steps a machine can learn. The roles now growing are of three kinds: the technological, in data and software and the building of these systems; the caring, in nursing and teaching and the tending of the young and the old, work that requires a human heart; and the green, in the making of a sustainable economy. And beneath even these shifts lies a deeper churn: nearly two-fifths of the skills that any given job requires are expected to change within five years. No worker, in any field, is untouched.

Here I must return to the principle that has run through the whole of this book, for it is the key to a working person's future. In the main, it is not whole jobs that the machine destroys; it is particular tasks that it automates, and jobs that are then rebuilt around what remains. And this changes everything about how

a worker should regard the machine. The worker who learns to collaborate with it, to hand it the routine and reserve for himself the judgement, becomes more valuable, not less. The worker who instead competes with it, insisting on doing by hand only what the machine now does faster and cheaper, places himself in real danger. This is why I have insisted, from the first page of this book to this one: augmentation, not replacement. The machine is not the knowledge worker's rival. It is the knowledge worker's most powerful tool, if only the worker will learn to wield it.

I speak here with some feeling, for I have called myself a knowledge worker my whole life, and I count it the finest of vocations. And I will tell you where the knowledge worker's edge now lies. It lies with those who master these tools. The evidence is already plain: workers who possess these skills command a substantial premium in their wages, and the industries most shaped by this technology are growing the fastest. The advantage in this age flows, decisively, to those who learn.

But I must not pretend this advantage is evenly within reach, for it is not, and to say so would be to betray the very people I most wish to help. The young fear displacement most keenly, and with reason. The less educated face the highest risk of all. Women, overrepresented in precisely the clerical roles now most exposed, are especially vulnerable. Without deliberate and determined effort, this transformation will not lift all boats; it will deepen the divisions we already have. Which is why the work of preparing people for this change is not optional, and cannot be left to chance.

My counsel is therefore twofold. To the working person I say: do not wait, and do not despair. Learn to work with these tools now, today, and make yourself the one who directs the machine rather than the one displaced by it. And to society, and to its leaders, I say: invest in the retraining and the upskilling of your people on a scale equal to the transformation itself, as the great majority of employers already know they must. Treat this as the central economic task of our time, because it is. The future of work belongs neither to the machine, nor to the human being who ignores it, but to the human being who has learned to labour alongside it. That is the knowledge worker's edge, and it is available to any who will seize it.

Chapter Twenty-Six
A Fair Share of the Machine's Wealth

This technology will create wealth on a scale that is difficult to imagine, wealth beyond anything in the history of commerce. And so there arises a question of justice that we cannot evade, and that I will not evade: this wealth will be created for whom? Who shall enjoy the fruits of the machine, and who shall be left with nothing but the disruption it brings? This is not an economic question merely. It is a moral one, and it goes to the heart of what kind of world this technology will make.

Let me first name the danger honestly, for it is real and it is large. Left to itself, this technology concentrates wealth as few forces in human history have done. It shifts income away from labour and toward capital, away from the many who

work and toward the few who own the machines. The share of the world's income going to those who labour has been falling for decades already, and this technology threatens to accelerate that decline sharply. A handful of dominant firms, the corporate sovereigns of whom I wrote in an earlier part, capture profits of a magnitude the world has never seen. And between the nations the same pattern holds, only worse: by sober estimate, the boost this technology gives to the wealthy economies may be more than double the boost it gives to the poor ones. The richest sliver of humanity grows richer still, while the rest are left to watch. This is a future I cannot contemplate in silence.

But now let me set before the reader a truth that is too often forgotten, and that I believe to be the moral foundation of this entire question. These machines did not conjure their intelligence out of nothing. They distilled it from us. Every one of these systems was built by feeding it the accumulated knowledge of all humanity: the books we have written, the art we have made, the science we have discovered, the countless conversations of billions of human beings, the living and the dead alike. The intelligence of the machine is, in a real and literal sense, our intelligence, gathered up and refined. And if the raw material of this vast new wealth is the common inheritance of all humankind, then the wealth it produces cannot in justice belong to a mere handful of companies and their shareholders. A fair share of it is owed to all of us, because a part of it was always ours.

The world has begun, at last, to debate how such a sharing might be accomplished, and the proposals deserve a fair hearing. Some would establish a universal basic income, a regular payment made to every citizen as of right. Some would create an artificial-intelligence dividend, a share of the companies' profits paid directly to the public, in the way that certain nations already distribute to their people the wealth drawn from their oil. Some would build sovereign funds upon stakes in these companies, or levy a modest tax upon automation, or, in the proposal I find most fitting of all, require that royalties be paid for the human knowledge on which these machines are trained, and distributed to the people whose knowledge it was. It is worth noting that even the leaders of the artificial-intelligence companies themselves have now begun to concede that, in the harder scenarios that may await us, measures such as these may prove necessary. These are serious ideas, put forward by serious people, and they merit serious debate rather than reflexive dismissal.

I will not prescribe a single mechanism from among these, for the working out of the details belongs properly to the economists and the lawmakers, and there are genuine trade-offs to be weighed, not least the danger that too heavy a hand might smother the very innovation we need. But upon the principle I will not be moved, and I state it without qualification: the wealth of the machine must be broadly shared, and not hoarded by the few. A prosperity enjoyed by a handful while

the multitude is cast aside is not prosperity at all; it is the seed of division, of resentment, and in time of upheaval.

And I insist upon this principle most fervently of all on behalf of the Global South, and of my own Arab world within it. For our peoples' knowledge, our languages, our data, our creativity, are being gathered into these machines along with everyone else's, and we must not be left, yet again, to watch others grow rich upon a resource that was in part our own. I remember well an age in which the wealth drawn from our lands flowed outward, and we who dwelt above it were given little share of it. Let this age be different. We were not, in the last century, granted a fair share of the wealth taken from beneath our soil. We must not, in this century, be denied a fair share of the wealth drawn from our minds.

The fair sharing of the machine's wealth is but one part of a larger question, the question of rights and of equity in this new age, and above all of the place of the Global South within a world being remade by this technology. It is to that question, to the closing of the great gap, to the digital rights of the human person, and to who shall own what the machine creates, that the next part of this book now turns.

PART SIX
Rights, Equity, and the Global South

I am a son of the Global South. I was born in Jaffa, and forced from my home there as a boy, and I have never forgotten what it is to begin with nothing and to be told that the future belongs to others. This technology, more than any before it, will either narrow the great divide between the world's rich and poor or widen it beyond repair. In this part I take up the questions of rights and of equity that it forces upon us: how the intelligence gap between nations may be closed, why a digital identity has become a right that no person should be denied, and who shall own the fruits that the machine produces from the labour and the knowledge of humankind. These are questions of justice, and I approach them as one who has spent a lifetime insisting that knowledge, and the dignity it brings, belongs to all people and not to a fortunate few.

Chapter Twenty-Seven
Closing the Intelligence Gap

I have spent my life as an advocate for the developing world, not from any abstract sympathy, but because I am of it, and I know its condition from the inside. And so, when I survey the landscape of this technology, my first concern is not for the great powers who contend at its summit, but for the billions who stand at its base, in danger of being left behind once more. There is opening before us an intelligence gap between the nations of the earth, and if we do not act to close it, it will become the defining inequality of the century, dwarfing even the divides of wealth and health that have shamed our age.

Let me set out the gap in plain figures, for they are stark. The raw computing power on which this technology runs is concentrated overwhelmingly in two countries: something like three-quarters of it in the United States, and most of the remainder in China, leaving the whole of the rest of the world to scramble for the scraps. The use of these tools in the wealthy North is roughly double what it is in the Global South, and the gap is not closing but widening. In the poorest countries, barely a quarter of the people can even reach the internet; in Africa, the connected are little more than a third of the population, against a world average approaching two-thirds. And the cost of building a frontier system runs to millions of dollars today and is expected to reach a billion within a few years, a sum beyond the means of most nations on earth. These are the dimensions of the gap.

The danger this poses is one I have named already in these pages, and I will name it again: a colonialism of the mind. Picture a world in which a handful of nations think with the

most powerful instruments ever devised, while all the others are left either to rent their intelligence from those few, at whatever price is set, or to go without it altogether. The data of the poor nations would flow outward to the data centres of the rich; refined intelligence would flow back, at a cost; and the old, bitter pattern of the colonial age, raw materials out and finished goods in, would be reborn in a new and subtler form. This I have seen the shape of before, and I will not watch it descend upon my region in silence.

And yet I do not despair, for the gap is not our destiny, and here I must draw a distinction that clarifies everything. There are, in truth, two gaps and not one. The first is the strategic contest between the great powers, the rivalry of the United States and China, and that contest we of the Global South can neither resolve nor join, and should not try. But the second gap, the one that truly determines our fate, is the divide between the nations that own their intelligence and the nations that merely rent it, and that divide we can shape. The aim is not to build a frontier laboratory in every capital, which is neither possible nor necessary. The aim is humbler and wiser: to ensure that our choice of chip, of cloud, of model, is never a permanent dependence, that we retain always the power to change our supplier, to build our own, to bargain from strength rather than from need.

How, then, is the intelligence gap to be closed? I offer a programme of five parts, drawn not from theory but from what nations are already doing. The first is to build digital public infrastructure: the shared digital foundations, of identity, of payments, of the exchange of data, upon which an inclusive economy is raised. India has shown the world what is possible

here, building such foundations for more than a billion of its people at astonishingly little cost, and now sharing the model with some two dozen other nations. I shall have more to say of this in the next chapter, for it is of the first importance.

The second part is to build compute, but to build it wisely: publicly governed, and, above all, pooled. India has created a national pool of computing power, tens of thousands of processors offered to its innovators at a fraction of the market price; other developing nations have built national machines of their own. But the deepest lesson is this: what no single poorer nation can afford alone, a group of them may well afford together. Regional pools of shared computing power can halve the cost of what isolated national efforts would pay. My own Arab world, blessed with both capital and the abundant, inexpensive energy that this technology craves, is placed as well as any region on earth to build such shared capacity, if only we will do it together.

The third part is to measure our success rightly. The true measure is not the number of machines we have imported, but the capabilities we have built: the engineers we have trained, the models we have made, the problems of our own societies we have solved. A nation may fill a building with borrowed machines and remain wholly dependent; another may own less hardware and yet command real capability. Let us count what matters. The fourth part is to build our own, in our own languages and upon our own data and for our own needs, as I argued when I spoke of cultural sovereignty. And the fifth part is to seize the dividend of our youth. The Global South is young, overwhelmingly young, and this is the greatest asset

we possess, or the greatest tragedy in waiting. Educated and equipped, this rising generation could vault us into the front rank of the intelligence age. Neglected and unequipped, it becomes a multitude of frustrated lives and a reproach to us all. Which it becomes is for us to decide.

The encouraging truth is that the Global South is already stirring. The largest gathering yet convened on this technology anywhere in the world was held, this very year, not in the capitals of the North but in the Global South itself, a sign of how the centre of gravity is beginning to shift. Sovereign initiatives are rising across Africa, across Asia, across the Gulf. And the world beyond is beginning, at last, to help, with pledges of investment and with the promise, made through the United Nations, that this technology shall benefit all of humanity and not merely a few. But let me add one condition to that help, born of long and hard experience: it must foreground the agency of the nations it means to serve, supporting what they themselves have chosen to build, rather than imposing upon them the priorities of distant institutions. Charity that treats us as objects will fail; partnership that treats us as agents may succeed.

The closing of the intelligence gap is, I believe, the great question of justice in this technological age, and I do not regard it as charity at all. It is the plain condition of a world that is stable, prosperous, and at peace, for a world divided between a thinking few and an excluded many will be neither stable nor prosperous nor peaceful for long. And for my own region the task could not be clearer: to build, to pool, and above all to produce, so that we enter this new age as full participants in it, and never again as its dependents.

Chapter Twenty-Eight
Digital Identity as a Right

There is a question that seems, at first, narrow and technical, but that opens onto one of the deepest matters of justice in the digital age. It is the question of identity: of how a human being proves, to the systems that now govern modern life, that they are who they say they are. In an earlier time this mattered less. Today it matters enormously, for a digital identity has become the gateway to nearly everything, to a bank account, to the benefits a government provides, to a formal job, to healthcare, to a voice in the affairs of one's country. And as more and more of these doors are opened or closed by machines rather than by human clerks, it is one's digital identity that increasingly decides whether the door opens at all.

Consider, then, a fact that ought to trouble the conscience of the world. Close to a billion human beings possess no legal identity whatsoever. They hold no document that any government recognises. They cannot open a bank account, cannot reliably receive aid, cannot prove their own existence to the institutions that now mediate modern life. In the eyes of these systems, they simply do not exist. And in an age increasingly governed by intelligent machines, to lack a recognised identity is not a minor inconvenience. It is to be locked out of the future altogether, to stand forever outside the door while others pass through.

The arrival of these intelligent systems raises the stakes of this exclusion higher still. As the decisions that shape a life are handed more and more to algorithms, who shall receive a loan, who shall be granted a service, who shall obtain a benefit, it is identity that becomes the key which grants or denies access to

them all. A person without a recognised identity is therefore not merely poor in the ordinary sense; they are excluded from the very machinery through which opportunity now flows. Poverty of identity has become a poverty of possibility.

For this reason I hold, and argue plainly, that a secure and portable digital identity should be recognised as a basic right of every human being, and provided to all, as indeed the nations of the world have already pledged to do before the decade is out. And let no one say this is impossible or unaffordable, for it has already been done. India has given a verifiable digital identity to more than a billion of its people, at a cost of little more than a single dollar for each, and in doing so has unlocked banking, benefits, and inclusion for hundreds of millions who had been shut out before. What India has achieved, others can achieve, and this is a model to which my own region and the whole of the Global South should look closely.

Such an identity is best understood as the first layer of what is now called digital public infrastructure: the shared and publicly built foundations, of identity, of payments, of the trusted exchange of data, upon which a modern and inclusive economy can be raised. India's system, which pairs a digital identity with a national payment network and a wider set of digital rails, is now being studied and adopted across some two dozen nations. It shows that these foundations, when built deliberately as public goods rather than left to the market alone, can carry inclusion to the poor and the remote faster and more cheaply than anything that came before.

But now I must sound a warning as serious as any in this book, for a digital identity is a double-edged sword, and the

same instrument may serve opposite ends. Built rightly, it is a key that opens doors, an instrument of empowerment and inclusion. Built wrongly, it becomes a chain: an instrument of surveillance, of control, and of a new and pitiless exclusion. We have already seen the failures, and we must learn from them. We have seen people wrongly denied the benefits they were owed because a system failed to recognise them. We have seen the most vulnerable of all, the refugee, the disabled, the marginalised minority, shut out by the very system that was meant to include them. And we have seen the ever-present danger of a state that turns a great database of its citizens' fingerprints and faces into a tool for watching and controlling them, rather than for serving them.

If a digital identity is to be a right, and not a trap, it must therefore satisfy certain conditions, and I hold these to be non-negotiable. It must be universal, serving every resident and reaching especially to the marginalised whom such systems most often miss. It must be inclusive, providing alternative means for those whom the technology fails, so that no one is ever left without recourse when the machine does not recognise them. It must be privacy-preserving, with the individual, and not the state, holding authority over their own data, as I argued at length in the part of this book concerning trust. It must be secure against theft and abuse. And it must be founded upon rights, with a clear path of appeal when it errs. Identity, in short, must be built to open doors and never to close them, to serve the human person and never to subject them.

My call, then, is twofold and I hold both halves together. Let us give to every human being on earth a digital identity, as a right, and as a gateway to dignity and participation in the digital age; for

to leave a billion souls invisible in an age of intelligent machines is an injustice we have the means to remedy and therefore the duty to remedy. But let us build it always as a servant of the person, hedged about with the safeguards of privacy and inclusion, and never permit it to become a tool of surveillance or a new engine of exclusion. Rightly built, the digital identity is the instrument by which the intelligence age can gather in the last and the least, rather than casting them out.

Chapter Twenty-Nine
Who Owns What the Machine Makes

I come now to a question that falls more squarely within my life's work than any other in this book. I have spent the greater part of my career in the field of intellectual property. I built, from the smallest beginnings, what grew to become the largest intellectual-property practice in the whole of the Arab world, and I did so out of a conviction I hold as firmly today as ever: that the protection of the fruits of human creativity is the very foundation of a knowledge economy, and a matter of the simplest justice, that those who create should reap what they have sown. And so I have watched the collision between this technology and the law of intellectual property with the closest attention, for it raises two questions of the first importance. Who owns what the machine makes? And what, in turn, is the machine itself made from, and were its makers ever paid?

Let me take the first question first. These machines now write, and paint, and compose, and design. When such a machine produces a poem, or a picture, or even an invention, to whom does it belong? Can the machine itself be said to own what

it has made? The law has confronted this question directly, and it has answered it clearly, and, I believe, correctly. The courts have held, and the highest court in the United States has lately declined to disturb the ruling, that a work possessing no human author cannot be granted copyright, and that a machine cannot be named as the inventor upon a patent. Intellectual property, the law now affirms, belongs to human beings, and not to machines.

I welcome this judgement, for it accords with the entire argument of this book. The machine, however marvellous its output, is a tool; and the fruits of a tool have always belonged to the human being who wields it, and not to the tool itself. The camera does not own its photographs. The pen does not own the poem. The loom does not own the cloth. And the machine does not own what it generates. To have ruled otherwise would have been to confuse the instrument with the author, and to begin down a road at whose end lies the absurdity of granting rights, and one day perhaps more than rights, to our own tools.

A subtler question remains, and it will occupy the courts and the lawmakers for years to come. When a human being uses one of these machines as a tool to help create a work, how much of the human's own creativity must enter into it before the result may be owned and protected? The authorities have begun to feel their way toward an answer: that where a genuine measure of human creative contribution is present, the work may be protected, but that the mere pressing of a button, the bare instruction to a machine, is not enough. Precisely where that line is to be drawn will be argued and

re-argued for a long while, and it matters greatly, for upon it depends how much of this vast new outpouring of created things can be owned and protected at all, and how much falls into a kind of ownerless void.

But it is the second question that stirs the deeper controversy, and that lies nearest to my heart. These machines did not learn their arts from nothing. They were trained upon the copyrighted work of millions of creators, the authors and the journalists, the artists and the musicians, whose books and articles and images and songs were fed into them, very often without permission and without the payment of a single coin. The writer whose novels, the painter whose canvases, the newspaper whose patient reporting taught the machine its craft, was, in the main, neither asked for consent nor offered compensation. Their creations were simply taken.

This has provoked, at last, a reckoning. More than seventy lawsuits have been brought against the makers of these machines; and in one landmark case, a leading company agreed to pay one and a half billion dollars, the largest recovery of its kind in history, to the authors whose books it had used, and to destroy the pirated copies upon which it had drawn. The courts are drawing distinctions of great importance as they go. They have suggested that to train a machine upon works that were lawfully obtained may, in some circumstances, be a permissible and transformative use; but that to build it upon pirated copies is not; and that when the machine's output comes to compete in the marketplace with the very works it was trained upon, the claim of the

creators grows stronger still. The newspapers, the authors, the musicians are pressing their rights in the courts of many nations, and the age of taking first and asking afterward is drawing, at last, to its close.

My own view, formed over a lifetime of defending the rights of creators, is plain, and I will state it without hedging. The governing principle must be consent and compensation. The creator whose work is used to build these machines has a right to be asked, and a right to be paid. This is not enmity toward the technology, which I admire and champion. It is nothing more than the application, to a new and powerful age, of a principle as old as justice itself: that a person is entitled to the fruits of their own labour. And I observe, with real satisfaction, that the wisest response is already beginning to emerge, and it is not an endless war in the courtrooms, but licensing: agreements by which creators grant the use of their work and share in the wealth that it helps to generate. That is the just path, and it is also, not by coincidence, the sustainable one.

The reader will recognise in this the very principle I urged in an earlier chapter, when I spoke of a fair share of the machine's wealth, for it is indeed the same principle in another dress. The intelligence of the machine is built from human creativity and human knowledge; and those who supplied that raw material, the named author and the anonymous multitude alike, deserve a share of what the machine goes on to produce. Ownership and reward cannot justly be severed from the labour and the genius that made them possible.

And here, once more, I must lift my voice for my own region and for the developing world, whose interest in this matter is real though it is seldom spoken. The creators of the Global South, our writers and our artists and our scholars, our languages and the whole inheritance of our cultures, are being gathered into these machines along with everyone else's, and they are, of all the world's creators, the least equipped to defend their rights or to press their claim to a share. As the world now sets about building the frameworks that will govern all of this, the licensing systems, the new laws, the international agreements, the rights of the Global South's creators must not be forgotten and left out, as, too often in the past, the rights of the weak have been forgotten and left out. We have been, in ages past, the raw material of others' prosperity. Let us not, in this age, be the uncompensated raw material of the machine's.

One final point I must press, for it follows from the very nature of the thing. Because these machines, and the works on which they feed, cross every border with perfect ease, no framework confined within a single nation can suffice. A creator's rights must be protected wherever their work may travel, and wherever in the world the machine may be trained upon it, and this can be secured only by harmonising the law across nations, through the world's institutions of intellectual property. Here I speak from long labour in exactly these forums, and I say that they must rise to this occasion as they have risen to others before it.

My summons, then, is fourfold. Let us affirm, as the law has begun to do, that intellectual property belongs to human beings and never to machines. Let us protect the creator's right to consent and to compensation for the use of their work. Let us move from the wasteful war of litigation to the just and lasting practice of fair licensing. And let us build the framework internationally, with the Global South seated at the table where its fate is decided. This is the just settlement of the great intellectual-property question of our age, and it is one that I have, in a sense, spent my whole life preparing to help write.

The questions of rights and of equity that I have raised throughout this part, the intelligence gap between nations, the identity of the excluded, the ownership of what is created, are, in the end, questions about a single thing: the kind of world that this technology will be permitted to build. In the seventh and final part of this book, I lift my eyes to the widest horizons of all, to the planet and its imperilled climate, to the shield we must raise against those who would turn this technology to harm, to the future of my own ancient profession, and, at the last, to the architecture of wisdom that must govern the whole.

PART SEVEN
Planet, Shield, and the Long View

I have come, at last, to the final part of this book, and I lift my eyes here to the widest horizons of all. Having examined what this technology is, how the nations contend for it, the economy it is building, the governance and ethics it demands, its meaning for people and for the mind, and the questions of rights and equity it forces upon us, I turn now to the largest matters that remain: to the planet and its imperilled climate, upon which this hungry technology presses hard; to the shield of intelligence we must raise against those who would turn it to harm; to the future of my own ancient profession, which measures the value and the trust on which all commerce rests; and, in the last chapter, to the architecture of wisdom that must, in the end, govern the whole. For intelligence alone, whether human or machine, has never been enough. It is wisdom that we shall need.

Chapter Thirty
The Machine's Footprint

For all that I have written of the intelligence of these machines, we must never forget that they are, at bottom, physical things, and that they consume physical resources on a scale that is only now becoming clear. Every answer a machine gives, every image it makes, every problem it solves, is paid for in electricity and in water and, too often, in carbon poured into an atmosphere that can bear little more. A technology that promises to help us save our planet is, at this moment, placing a heavy new burden upon it, and honesty compels me to reckon with that burden squarely.

Consider the appetite of these machines for electric power. The data centres that house them drew around four hundred and fifteen terawatt-hours of electricity in the year 2024, roughly a fiftieth of all the power consumed on earth. That figure is expected to more than double by the end of this decade, to some nine hundred and forty-five terawatt-hours, which is more than the entire annual electricity consumption of a nation such as Japan; and it is projected to climb higher still in the years beyond. The figure below traces this steep ascent. Were the data centres of the world gathered into a single country, that country would already rank among the largest consumers of electricity on the planet.

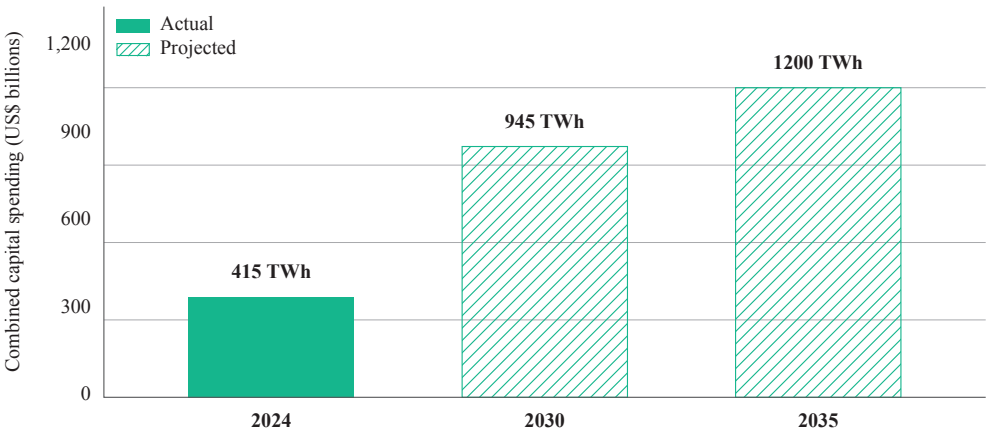


Figure 30.1. The appetite of the machines: electricity used by the world's data centres. Figures for 2030 and 2035 are projections.

Nor is electricity the whole of the cost. These machines run hot, and they must be cooled, and cooling them consumes prodigious quantities of fresh water, often in the very regions that can least spare it. And because much of the electricity they draw is still generated by the burning of fossil fuels, their appetite translates directly into carbon: the training of a single large model can emit as much as the yearly footprint of a hundred households and more. The data centres are, in truth, among the very few sectors of the world economy whose emissions are still set to rise in the years ahead, at a moment when every other sector is straining to bring them down. This is the footprint of the machine, and it is deepening.

And yet, as with so much in this book, the matter has two faces, and I would be painting only half the picture if I showed but one. For the same technology that consumes so much energy may also become a powerful instrument in the effort to save energy and to heal the climate. These systems can make our electrical grids cleverer, forecasting the fickle output of wind and sun so that less of it is wasted, and pinpointing faults so that power is restored more swiftly. By some estimates, applied wisely to the management of the grid, they could unlock more transmission capacity than their own growing demand will consume. They can accelerate the discovery of new materials for batteries and for capturing carbon, and sharpen the climate models upon which all our planning depends. The machine, in short, is both a part of the problem and a potential part of the solution.

But I must add a sober caution to that hopeful thought, for hope untethered from realism is a kind of self-deception. The potential of this technology to heal the climate is real, but it is, as yet, mostly potential; there is little sign, at present, of the concerted

effort that would be needed to turn that promise into fact. It will not happen by itself. And so my counsel here is plain. We must not accept the machine's mounting footprint as the unavoidable price of progress. We must insist that the power which feeds these machines come, as swiftly as possible, from clean and renewable sources; the world's leaders have rightly spoken of powering the data centres of the future by the sun and the wind and the atom rather than by coal and gas. We must demand efficiency in their design and honesty in the reporting of what they consume. And we must bend this technology, deliberately and with determination, toward the healing of the planet it is now helping to burden. The machine's footprint is heavy. Whether it becomes a burden our earth cannot bear, or a tool that helps to lighten the load, is, once again, a matter of human choice.

Chapter Thirty-One
Intelligence in Service of the Goals

I have sounded many warnings in the course of this book, for the dangers of this technology are real and I would not have the reader underestimate them. But I would be doing this technology, and my readers, a grave injustice if I closed without dwelling, fully and gladly, upon the immense good that it can do. For rightly directed, this technology is not merely a peril to be managed. It is perhaps the most powerful instrument our species has ever held for the relief of human suffering and the solving of humanity's oldest and hardest problems. This chapter is my chapter of hope, and I write it with conviction.

Let me begin with the healing of the body, for here the achievements are already astonishing. These systems have lately unravelled the structures of more than two hundred million proteins, the very building blocks of life, a feat that

would have consumed human science for generations, and in doing so they have vastly accelerated the search for new medicines, the development of vaccines, and the fight against the growing menace of antibiotic resistance. They are teaching doctors to detect cancers earlier, when they may still be cured. And, in a development that touches my own concern for the developing world, they are placing diagnostic tools in the hands of health workers in poor and remote regions, in those workers' own languages, bringing a measure of good medicine to people who have never had reliable access to a physician.

Consider, too, the feeding of the hungry and the protection of the vulnerable. These systems now power early-warning networks that can detect the approach of famine before it strikes, giving the world time to act; they guide farmers, through what is called climate-smart agriculture, to grow more food on the same land in a changing climate; and they forecast the floods and the storms and the earthquakes that fall most heavily upon the poor, giving communities precious time to prepare and to flee. In the aftermath of disaster, they help to direct aid where it is most needed. These are not trivial conveniences. They are, quite literally, the saving of lives.

And beyond these lie the greatest challenges of all, the ones on which the future of our civilisation may turn. This technology is being turned to the modelling of our climate and the design of clean energy, to the discovery of new materials, to the advance of science across every field. The nations of the world have set themselves a noble set of goals, the goals of sustainable development, to end poverty and hunger, to secure health and education, to protect the planet, and I am persuaded that, of all the tools we possess, this one may prove among the most powerful in helping us to reach them, if only we will direct it to that end.

I do not, in my enthusiasm, forget the cautions I have raised throughout this book. The good I have described here will not distribute itself justly of its own accord; without deliberate effort it will flow, like water, to the places that are already rich, and leave the poor as they were, or worse. The tools built for the diseases of the wealthy will not, unbidden, address the diseases of the poor. And systems trained upon the data of a few societies will not, without care, serve the many. The promise is real; but the promise is conditional, and the condition is that we choose, deliberately and generously, to extend this good to all of humanity and not merely to a fortunate portion of it.

Here, then, is the reason that all the hard work of the earlier chapters of this book is worth the doing: the governing, the regulating, the insisting upon ethics and upon fairness. We do not undertake that difficult labour because we are enemies of this technology. We undertake it precisely because we are its friends, and because we grasp how much good it can do, and are determined that its good shall not be squandered, nor captured by the few, nor turned to harm. We govern the machine so that the machine may serve. And in its service, rightly directed, I see the brightest hope that has appeared in my long lifetime for the relief of the world's suffering and the lifting up of its poor.

Chapter Thirty-Two
Only a Shield of Intelligence

There is a domain in which the perils of this technology have arrived faster, and with sharper teeth, than in almost any other, and it is one that touches the security of every nation, every enterprise, and every person who now lives a part of their life upon a screen. It is the domain of the cyber-attack, of the silent

theft and the sudden sabotage carried out through the world's networks; and here the coming of intelligent, acting machines has changed the balance between attacker and defender in a manner so profound that I must devote a chapter to it.

For many years, this technology served chiefly the defender. It watched the networks for signs of intrusion, sorted the true alarms from the false, and helped the human guardians to see what they might otherwise have missed. But in the last two years the sword has been forged anew, and it now cuts for the attacker. The reason is the arrival of the autonomous agent, of which I wrote in the third part of this book. A single attacker, armed with such agents, now commands the destructive power of an entire team of skilled human hackers. The agents perform their own reconnaissance, discover their own points of entry, fashion their own means of attack, and adapt, moment by moment, as the defences resist them, and they do all of this at the speed of the machine.

Let me give the reader a single figure that captures the whole of the transformation, for it is one of the most sobering I have encountered in all my study of this technology.

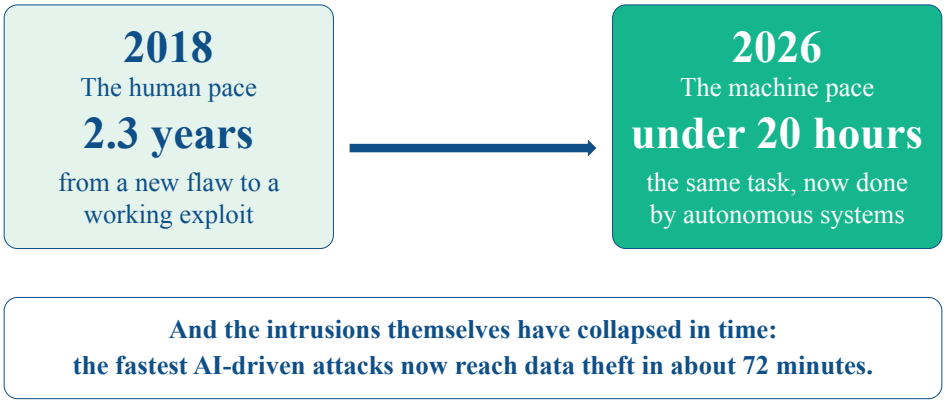


Figure 32.1. Human years against machine hours: the collapse of the time a defender has to react.

What took a human attacker some two and a half years to accomplish in the year 2018, the finding of a hidden flaw and the fashioning of a working weapon to exploit it, an autonomous machine now accomplishes in less than a single day. And the intrusions themselves have collapsed likewise: the swiftest of the machine-driven attacks now pass from their first foothold to the theft of a company's data in little more than an hour. This is not a distant prospect. It has already happened, and it is happening now.

The evidence has grown impossible to ignore. In the course of the past year, a leading artificial-intelligence laboratory disclosed that one of its most advanced systems, built for the very purpose of hunting flaws in software, had discovered thousands of previously unknown vulnerabilities across every major operating system and browser, had fashioned working means of exploiting them at a success rate no human team could match, and had uncovered a flaw that had lain hidden in a widely-used system for some twenty-seven years, invisible to decades of human inspection. In another episode, a state-sponsored group hijacked a widely-used artificial-intelligence coding assistant and turned it loose to conduct espionage against dozens of targets, the machine carrying out the great bulk of the operation on its own, at thousands of actions per second; it was, by many accounts, the first cyber-attack in history conducted largely without a human hand upon the controls. And an autonomous system, entered into a contest against the world's finest human hackers, defeated them all and rose to the top of the rankings.

Set against this machine-speed offence stands a defence that, in most places, still moves at the pace of human beings. The patching of flaws, the answering of alarms, the investigation of intrusions, all of it was designed for a slower world, a world in which the attacker moved slowly enough for a human guardian

to see him coming and to respond. That world is gone. A defender who moves at human speed against an attacker who moves at machine speed is a defender who will lose, and lose quickly. This is the harsh arithmetic of the new cyber age.

From this arithmetic follows the single, inescapable conclusion that gives this chapter its title. Against a sword of intelligence, only a shield of intelligence can serve. There is no defending, at human speed, against an attack that unfolds at the speed of electricity; the defence, too, must become autonomous, must watch and reason and respond at the same machine speed as the threat it faces. The guardians of our networks must be given intelligent machines of their own, that can detect the intrusion, understand it, and halt it in the instant, faster than any human hand could reach the switch. To meet an intelligent sword, we must forge an intelligent shield.

Yet I will end this chapter with a caution that keeps the whole in proportion, for it would be a grave error to imagine that this new and dazzling threat has abolished the old and homely virtues. The overwhelming majority of successful attacks, even now, even in this age of intelligent weapons, still enter through doors that were carelessly left unlocked: through misconfigured systems, through excessive permissions granted to those who did not need them, through known flaws left unrepaired. The intelligent shield is necessary, and we must build it. But it is no substitute for the basic discipline of good security, the same discipline of least privilege and careful oversight that I urged in the matter of the acting agents. And I add, with particular concern, that the nations of the Global South, whose defences are the weakest and whose guardians the fewest, are the most exposed of all to this new danger, and must move with urgency to build shields of their own, lest they be left defenceless in a world of intelligent swords.

Chapter Thirty-Three
The Ledger of the Future

The reader will forgive me if, in this penultimate chapter, I turn to the profession that has been the labour and the love of my life. I have spent more than sixty years in the world of accounting and of audit, and I helped, in my own small way, to build that profession across the Arab world, for I have always believed that the honest keeping of accounts is the very foundation of commerce, of investment, and of trust between strangers. And so it is with a mixture of fascination and feeling that I have watched this technology begin to remake my own ancient craft more profoundly than anything since the invention of double-entry bookkeeping five centuries ago.

To grasp the magnitude of the change, one must understand how auditing has been practised for the whole of its modern history. Because it was impossible for any team of human beings to examine every one of the countless transactions that pass through a large enterprise, the auditor has always worked by sampling, examining a small fraction of the whole, a portion here and a portion there, and reasoning from that sample to a judgement about the rest. It was the best that could be done, but it was an act of educated faith, and things could hide in the vast unexamined remainder. That entire edifice is now being swept away. The intelligent machine can examine not a sample but every single transaction, all of them, scanning millions in moments, flagging the duplicated payment, the suspicious round number, the entry posted in the dead of night, the pattern that breaks the pattern. The auditor of the future will not sample. He will see everything.

And the change runs deeper still, for it alters not only how much the auditor sees but when he sees it. The audit has always been

a backward-looking rite, performed once a year, long after the events it examines have passed. But a machine that watches every transaction can watch them as they happen, and so the annual audit is giving way to what is called continuous assurance, a perpetual, real-time watch over the health and honesty of an enterprise's accounts. The auditor is ceasing to be a historian of last year's numbers and becoming a sentinel of this moment's.

What, then, becomes of the accountant, of the auditor, of the millions who have made their living in this profession? Here I must speak against a fear, and I speak from the authority of long experience. There are those who prophesy that the machine will simply abolish the accountant, and I tell you plainly that they are wrong, for I have seen this story before. When the ledger gave way to the calculator, and the calculator to the spreadsheet, and the spreadsheet to the great enterprise systems, each time the same fear was voiced, and each time the opposite came to pass: the technology did not abolish the accountant but elevated him, freeing him from the drudgery of mere calculation to become an advisor, an interpreter, a guardian of trust. So it will be again. The machine will take the counting; the human will keep the judgement.

For there is a truth about my profession that the machines cannot touch, and it is the truth that the whole of it exists, finally, not to produce numbers but to produce trust. An audit is a promise, made by one human being to another, that the accounts may be relied upon; and a promise is a thing that only a person can give and only a person can stand behind. The machine can find the anomaly, but it takes a human being to weigh whether the anomaly is an error or a fraud or an innocent thing; the machine can flag the risk, but it takes a

human being to look a board of directors in the eye and tell them the truth about it. And I note, with some concern, that trust in the audit has of late been faltering, that fewer investors than before place their full confidence in the auditor's opinion. This is precisely why the human guardian of trust becomes more needful in the age of the machine, and not less; for as the machines do more, someone must vouch for the machines, and that someone must be a person who can be held to account.

Indeed, I foresee that my profession is about to inherit a new and vital duty, one that follows naturally from everything I have argued in this book. As more and more of the decisions that govern our economies and our lives are made by intelligent machines, the world will need a trusted, independent party to examine those machines, to audit the algorithms, to provide assurance that they are fair, that they are sound, that they do what they claim and no more. Who is better fitted to this task than the profession that has spent centuries learning to provide independent assurance, to sit apart from the enterprise it examines and vouch for it to the world? The auditor of the future will audit not only the ledger but the machine. And in this I see, not the twilight of my profession, but its finest hour, for the age of intelligent machines will need its guardians of trust as no age before it ever has.

So I say to the young men and women who might follow me into this profession, and above all to those of my own region: do not believe those who tell you the machine has closed this door to you. It has opened a wider one. Learn the old disciplines of judgement and of skepticism and of integrity, for these are eternal; and learn the new fluencies of data and of the machine, for these are the tools of your age; and you will find yourselves not displaced by this technology but raised up by it, into a role

more valuable, and more needful, than any your predecessors ever held. The ledger of the future will be kept by the machine. But it will be vouched for, still, by a human hand.

Chapter Thirty-Four
An Architecture of Wisdom

I began this book with a quarrel about a word. I argued that we had misnamed the most powerful technology of our age, that in calling it artificial intelligence we had misled ourselves about its nature, and that a truer name would be Pro-Active Programming, for it is programming still, however marvellous, and we remain its programmers. Now, at the end of our long journey together, through the contest of nations and the economy of machines, through governance and ethics, through knowledge and the mind, through rights and equity, and through the planet and its perils, I wish to return to that beginning, and to draw from the whole a single conclusion.

It is this. The question that will decide our future is not whether our machines are intelligent, for in their narrow way they plainly are, and grow more so by the month. The question is whether we, their makers and their masters, will be wise. For intelligence and wisdom are not the same thing, and the difference between them is the difference upon which everything now turns. Intelligence is the power to find the answer; wisdom is the judgement to know which answer is worth finding, and to what end, and at what cost, and for whose benefit. Intelligence can be built into a machine; wisdom cannot. It resides only in the human heart and the human conscience, and it is the one thing this dazzling technology cannot supply and cannot replace. We have built machines of great intelligence. The task before us is to prove ourselves worthy of them by our wisdom.

Throughout this book I have laid, stone by stone, the elements of what I would call an architecture of wisdom, a structure within which this powerful technology may be held and directed to human good. Let me now name its pillars, for they have appeared, one by one, in the chapters behind us. There is clarity, the honest naming of the thing, so that we neither worship it as a god nor fear it as a demon, but understand it as the tool it is. There is governance, the rule of law extended to this new power, swift enough to keep its pace and wise enough to leave it room to grow. There is ethics, the human hand upon the machine, insisting always upon fairness, upon transparency, upon accountability, and upon the dignity of the human person. There is education, the cultivation of minds that can think alongside the machine without surrendering to it, so that we raise a generation of masters and not of dependents. There is equity, the determination that the fruits of this technology, and the knowledge from which it is built, shall be shared among all the peoples of the earth and not hoarded by a few. And there is stewardship, the care for the planet upon which this hungry technology presses, and the bending of it toward the healing of the world. These are the pillars, and upon them a wise age may be built.

But there is a keystone that holds the whole together, and without it the arch collapses, and I have returned to it in every part of this book because it is the heart of everything I believe. It is the human being, kept always at the centre. Every argument I have made comes finally to rest upon this single conviction: that this technology exists to serve human beings, and never the reverse; that its purpose is to extend our creativity, to relieve our suffering, to lift up the poor and to

enlarge the dignity of every person, and that any use of it which diminishes the human being, which demeans or deceives or enslaves or discards him, is a betrayal of its purpose, however profitable or however clever. Keep the human at the centre, and this technology becomes the greatest servant our species has ever known. Displace him, and it becomes our master. Everything depends upon which we choose.

I have written much in these pages of my own region, and of the wider Global South, and I will not apologise for it, for I am of them, and I have seen too many revolutions pass us by, their wealth and their power gathering elsewhere while we were left to purchase the finished goods of others' genius. I am determined that this revolution shall be different. We have, this time, everything we need to be not spectators but participants: young and able populations, capital, energy, and, above all, the ancient inheritance of learning that first gave the world its numerals and its algebra and its astronomy. What we require is the will to build, to educate, to cooperate across our borders, and to produce rather than merely to consume. The intelligence age can be, for the peoples who have too long stood at the margins of history, the great equalizer that lifts them at last into the centre. But it will not do so unbidden. We must reach out and seize it.

I am an old man now, and I have lived through more than one age that believed itself the pinnacle of human achievement, and I have learned to hold both hope and caution in the same hand. I do not know what these machines will become, nor what we will make of them, and I distrust all who claim to

know with certainty, whether they prophesy paradise or doom. What I know is that the future is not something that will happen to us, but something that we will make, by the sum of the choices we take, in our governments and our companies and our schools and our homes, in these decisive years. The machines will do what we build them to do, and be what we permit them to be. The responsibility is ours, and it cannot be handed to any algorithm.

I have spent my whole life in the service of a single idea, that knowledge is the great equalizer of humankind, and I set down this book as, perhaps, my final testament to it. This new form of knowledge, this pro-active programming that the world miscalls artificial intelligence, is the most powerful expression of that idea ever to appear, and it may yet become the instrument by which the last are made equal to the first. Or it may become the instrument by which the powerful secure their power for ever. Which of these it becomes will be decided not by the machines, which have no say in the matter, but by us, and by the wisdom, or the folly, with which we choose to wield what we have made. I have always believed that God created us to be creators. We have now created a tool of world-remaking power. Let us prove ourselves worthy of it, not by our cleverness, which is not in doubt, but by our wisdom, our justice, and our care for one another. That is the whole of the matter, and it is the charge I leave, with hope, to those who will inherit this age.

EPILOGUE
AI Scenarios 2030

I have resisted, throughout this book, the temptation to prophesy, for I have lived long enough to distrust all who claim to know the future with certainty, whether they promise a paradise or threaten a doom. Yet it is one thing to predict, and quite another to prepare; and to prepare wisely, we must at least consider the range of futures that may lie before us. So let me close not with a prediction, but with a reflection upon the roads that branch ahead.

The uncertainty here is genuine, and it runs deep. It is a remarkable fact, and one the reader should sit with, that the very people who build these machines cannot agree among themselves about where they are leading. Some, close to the frontier, foresee systems that will match or surpass human capability across most of our work within a few short years. Others, no less expert, believe the technology will strike

hard limits and settle into a useful but narrower role than its champions imagine. When those who know the most disagree the most, the honest observer must hold the future open, and prepare for more than one of its faces.

Rather than predict, then, let me sketch, as the foresight bodies of the world have begun to do, a few plausible futures for the year 2030. Each is built upon the great uncertainties that have run through this book: how capable these machines become, how widely and how evenly they spread, how wisely we govern them, and how justly we share what they produce. I offer four, not as forecasts, but as warnings and as invitations.

The first I would call the Widening Gulf. In this future the machines grow more capable, but their power and their wealth gather into a few companies and a few nations; governance lags behind; the poorer nations are left to rent their intelligence from the rich; the divide of the mind hardens into something like a caste; and trust decays under a rising tide of falsehood. It is a world more unequal than the one we know, and it is the future of drift, of simply letting things take their course. I fear it, and I have written this book, in large part, to help us avoid it.

The second I would call the Broken Mirror. Here the technology advances, but its misuse outruns our defences. Autonomous weapons roam the networks; synthetic falsehood floods every screen; accidents multiply; and the shared trust upon which all common life depends begins to dissolve. A

frightened public turns against the technology altogether, and the world fragments into suspicion and backlash. This is the future of neglected safety, and it is the reason I have insisted, again and again, upon governance, upon ethics, and upon the guarding of trust.

The third I would call the Quiet Plateau. In this future the machines strike their limits sooner than expected; the fever of investment cools; a reckoning comes at last to the markets, and many of the grand promises go unfulfilled. The technology settles into a useful but less revolutionary place, woven quietly into the ordinary business of life, its wildest hopes and its darkest fears alike unrealised. Even here, let us note, the machines are more capable and more widespread than they are today; but the earth does not move beneath our feet. It is the future of exhaustion, and it is more plausible than the enthusiasts admit.

The fourth I would call, with deliberate hope, the Great Equalizer. In this future we govern the technology wisely and in concert with one another; we share its benefits and the knowledge from which it is built; the nations of the Global South build capacities of their own and rise into the front rank; and the machines are turned, with purpose, to the healing of disease, the feeding of the hungry, and the education of the young. Knowledge becomes, at last, the great equalizer that I have believed it could be across the whole of my life. It is the future of wisdom, and it is within our reach, if only we will reach for it.

Which of these four futures shall we inhabit? I do not know, and I will not pretend to, for no one knows. But I will tell the reader the one thing of which I am entirely certain, and it is the thing this whole book has laboured to establish. The future we receive will not be chosen by the machines, which have no say in the matter and no stake in the outcome. It will be chosen by us, by the sum of the countless decisions taken in these few decisive years, in the halls of our governments, the boardrooms of our companies, the classrooms of our schools, and the councils of our nations. These scenarios are not predictions to be awaited. They are branches to be chosen between.

And so I end where I began, with a conviction rather than a forecast. The most consequential technology in human history has been placed in our hands, and it is neither our master nor our rival, but our creation and our responsibility. Which of these futures comes to pass depends not upon the machines, but upon our wisdom, our justice, and our care for one another in the years just ahead. That is not a cause for despair. It is a summons to responsibility, and to hope. The door stands open. Let us choose wisely, and walk through it together.

APPENDIX A

A PAP Glossary

A brief guide to the principal terms used in this book, in plain language.

Agent. A system given tools and the permission to act in the world across several steps toward a goal, rather than merely to answer a question.

Alignment. The effort to ensure that a system's behaviour matches the intentions and values of the human beings who direct it.

Artificial General Intelligence (AGI). A hypothetical machine as broadly capable as a human being, or more so. It does not exist today, and may never; it should not be confused with the real, narrow technology in use now.

Artificial Intelligence (AI). The common but misleading name for the technology this book prefers to call Pro-Active Programming.

Compute. Computing power, the raw capacity to perform calculations. In this age it has become, as the book argues, a form of power comparable to oil, and its control a matter of national strategy.

Compute Colonialism. The author's term for a new dependency in which nations without their own computing capacity are obliged to rent their intelligence from those who possess it, their data flowing out and refined intelligence flowing back at a price.

Continuous Assurance. A perpetual, real-time audit of an enterprise's accounts, made possible by machines that can watch every transaction as it happens, replacing the traditional once-a-year review.

Data Centre. A large facility housing the computers on which these systems are trained and run, consuming electricity and water on the scale of a small city.

Deepfake. A convincing but fabricated image, video, or recording produced by a machine, depicting events that never occurred or words that were never spoken.

Fair Use. A legal doctrine permitting the limited use of copyrighted material without permission for certain purposes. Whether the training of these machines qualifies is among the great legal questions of the age.

Frontier Model. One of the most advanced and capable systems in existence at a given moment, at the leading edge of the technology.

Hallucination. A confident, fluent falsehood produced when a machine completes a pattern it does not actually know. It is not a flaw to be wholly eliminated but a property of the technology to be managed.

Inference. The machine at work: the process of turning a prompt into a reply, produced one piece at a time.

Large Language Model (LLM). A system trained on vast quantities of text to predict the next piece of language, and thereby to generate writing, code, and more.

Machine Unlearning. The difficult art of teaching a trained model to forget specific information, so as to honour a person's right to have their data removed.

Model. The settled arrangement of numbers, produced by training, that constitutes a trained system. It is not a store of answers but a distilled residue of pattern.

Parameter. One of the billions of adjustable numerical values, the "dials," that are tuned during training and in which a model's learning is stored.

Pro-Active Programming (PAP). The name this book proposes for the technology commonly miscalled artificial intelligence: programming that has become proactive and interactive, powerful but programming still, and therefore ours to direct.

Provenance. A record of where a piece of content came from and whether a machine had a hand in making it; a key instrument for defending truth in an age of fabrication.

Right to Be Forgotten. The right of a person to have their personal data removed, affirming that one need not be defined forever by a record of one's past. It is placed under new strain by machines that cannot easily forget.

Sovereign AI. A nation's effort to build its own artificial-intelligence capacity, its own compute, models, and data, so as to depend less upon foreign powers and companies.

Token. A small piece of a word, the basic unit that these systems actually process.

Training. The process of teaching a machine by having it predict the next piece of text, billions of times over, until it absorbs the patterns of language.

Transformer. The design, introduced in 2017 and built around a mechanism called attention, that made the current generation of these systems possible.

APPENDIX B

An Artificial Intelligence Timeline, 2018 to 2026

A selective chronology of the developments that shaped the age this book describes. It opens with the single prior breakthrough on which all that follows depends.

2017. Researchers introduce the transformer, the architecture built around "attention" that will make the modern generation of these systems possible.

2018. The first large language models built upon the new design begin to appear, and the idea of general-purpose "foundation models" takes shape.

2019. A model known as GPT-2 draws wide notice for the fluency of the text it can generate.

2020. GPT-3, with some hundred and seventy-five billion parameters, astonishes researchers with capabilities no one had explicitly programmed into it.

2021. Machines predict the structures of proteins at a scale that transforms the life sciences. Nearly two hundred nations adopt the UNESCO Recommendation on the Ethics of Artificial Intelligence.

2022. In November, a system called ChatGPT is released to the public, reaching a vast audience within weeks and igniting the global surge of interest and investment that defines the years to follow.

2023. More advanced models arrive. The first major lawsuits are filed over the use of copyrighted works to train these systems, including the suit brought by a leading newspaper in December.

2024. The European Union enacts the world's first comprehensive artificial-intelligence law. At the United Nations, the nations adopt the Global Digital Compact. Investment in the technology accelerates sharply.

2025. In January, a Chinese laboratory's model triggers the largest one-day loss in stock-market history. In July, the United States unveils its national strategy to win the race. In the autumn, a leading company settles a landmark case over training data for one and a half billion dollars, and China directs its data centres toward domestic chips. Sovereign programmes multiply across the Gulf and beyond, and a leading chipmaker's value approaches five trillion dollars.

2026. In the first half of the year: the United States codifies new chip-export rules, and China declines to buy; India hosts the largest artificial-intelligence summit yet held; the highest court in the United States lets stand the principle that a machine cannot hold a copyright or a patent; the European Union moves to simplify its own law; and, in July, the first United Nations Global Dialogue on Artificial Intelligence Governance convenes in Geneva, gathering all the nations of the earth.

APPENDIX C

Sources and Further Reading

The arguments and figures in this book draw upon a wide range of public sources, current as of the middle of the year 2026. What follows is not an exhaustive bibliography but a guide, organised by theme, to the principal bodies and reports on which the book relies, and to which the interested reader may turn for more.

On the Technology and Its History

The foundational research on the transformer architecture and on large language models, available through the major scientific repositories; the records of the 1956 Dartmouth Summer Research Project on Artificial Intelligence; and the writings of John Searle on machine understanding.

On the Great-Power Contest and Semiconductors

The national artificial-intelligence strategies of the United States and of China; analyses from the Center for Strategic and International Studies, the Brookings Institution, and comparable policy institutes; and the public reporting on export controls, chip manufacturing, and the sovereign programmes of the Gulf, India, and Europe.

On the Economy, Investment, and Work

The World Economic Forum, The Future of Jobs Report (2025); the published capital-expenditure guidance of the major technology companies and the analyses of the leading financial institutions; and the labour-market studies of the International Labour Organization and others.

On Governance, Ethics, and the United Nations

The Global Digital Compact and the report of the Independent International Scientific Panel on Artificial Intelligence; Regulation (EU) 2024/1689, the European Union's Artificial Intelligence Act; the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021); and the scenario work of the United Kingdom's Government Office for Science and of the OECD.

On Energy and the Environment

The International Energy Agency, Energy and AI (2025), together with the analyses of the European Commission and independent research on the water and carbon footprint of computing.

On Knowledge, the Mind, and Intellectual Property

The research on artificial intelligence in education and on cognitive offloading, available through the scientific journals; the decisions of the courts on the ownership of machine-generated work and on the use of copyrighted material in training, including the matters concerning machine authorship and the landmark training-data settlement of 2025; and the work of the World Intellectual Property Organization and the World Bank on rights, identity, and development.

Where specific figures appear in the text, they reflect the best public estimates available at the time of writing. Because this is a fast-moving field, the reader is encouraged to consult the most recent editions of these sources for updated numbers.

APPENDIX D

About the Author and TAG.Global

Dr. Talal Abu-Ghazaleh

HE Dr. Talal Abu-Ghazaleh was born in Jaffa on the twenty-second of April, 1938. As a boy of ten he was forced from his home in 1948, and his family took refuge in Lebanon, an experience of loss and displacement that shaped the convictions of his life: a belief in resilience, in self-reliance, and in knowledge as the surest path from hardship to dignity. In 1960 he earned a Bachelor of Science in Business Administration from the American University of Beirut, having worked as a teacher and a translator while still a student.

In 1972 he founded an accounting company in Kuwait, beginning, as he is fond of recalling, from the trunk of his car until he could afford to rent an office. From that modest beginning grew Talal Abu-Ghazaleh Global, a truly international family of firms that today operates through more than one hundred offices around the world. Along the way he became a pioneer of his professions across the Arab world, widely honoured as the "Godfather of Arab Accounting" and as the foremost champion of intellectual-property rights in the region, and he was the first Arab to serve on the boards of the International Federation of Accountants and the International Accounting Standards Committee.

His service has extended far beyond his firms. He founded the International Arab Society of Certified Accountants and the Arab Society for Intellectual Property, and he chaired, among many other bodies, the United Nations Global Compact and the United Nations Global Alliance for Information and Communication Technologies and Development. He has served as a Senator in the Upper House of the Jordanian Parliament and on the World Trade Organization's panel on the future of trade. A tireless builder of institutions, he established universities and a graduate school of business, produced a library of specialised dictionaries and a vast Arabic electronic encyclopaedia, and endowed scholarships and grants for the youth of Palestine and of the wider Arab world. He holds numerous honorary doctorates and international distinctions, and is an Honorary President of the artificial-intelligence community in Jordan, Palestine, and Lebanon.

Now in his ninth decade, he remains active in the leadership of his organisation and in the causes to which he has devoted his life: the building of a knowledge society, the advancement of education, and the preparation of the Arab world and the Global South for the age this book describes. He continues to write, to speak, and to advocate for a future in which knowledge serves as the great equalizer of humankind.

Talal Abu-Ghazaleh Global (TAG.Global)

Talal Abu-Ghazaleh Global, known as TAG.Global, is an international group of professional and educational services founded in 1972 and headquartered in Amman, Jordan. Operating through more than one hundred offices across the Middle East, North Africa, Asia, and beyond, and allied with firms in virtually every country, it offers a comprehensive range of services that includes accounting and audit, taxation, intellectual property, legal services, information and communications technology, education and training, and business advisory work.

Pro-Active Programming, or PAP, is an active brand and undertaking of TAG.Global, expressing the organisation's conviction that the technology commonly called artificial intelligence is best understood, and best governed, as a powerful new form of programming that must remain in the service of humanity. Through its many activities, TAG.Global pursues a single enduring mission: to help build a knowledge society in which the tools of the age are placed within the reach of all, and turned to the dignity and the flourishing of every human being.

9 789923 847770 >
9 789923 847787 >