



INVISIBLE SHIFT



Hakan Özgür

mindshiftme.de

INVISIBLE SHIFT

*Power, Machines, and the Question
of Who We Want to Be*

Hakan Özgür

MindShift Series — Book 6

www.mindshiftme.de

© 2026 Hakan Özgür

All rights reserved.

Email: yes@mindshiftme.de

As of: March 2026

For all who look closely.

Contents

PART I: THE NEW POWERS

1. The Invisible Revolution
2. The World Game — America, China, and the Void of Europe
3. The New Rulers — Tech Oligarchs and Power Without Mandate
4. The Economy of Abundance — and Why It Doesn't Work

PART II: THE DESTRUCTION OF REALITY

5. War Without a Battlefield
6. The Synthetic Public
7. Surveillance as a Business Model
8. Democracy in Intensive Care

PART III: THE SHAPING

9. The Return of the Political
10. The Architecture of Freedom
11. Being Human in the Age of Machines
12. Epilogue — Letter to a World in Transition

PART I

The New Powers

What is happening right now — and why it is unlike anything before

1

The Invisible Revolution

How a handful of companies shifted the architecture of global power

In November 2022, a company in San Francisco released a language model. It was called ChatGPT. Within five days, a million people had used it. Within two months, a hundred million. No product in the history of technology had spread faster.

That was three and a half years ago.

Since then, something has changed that is bigger than an app, bigger than an industry, bigger than a market. The architecture of power has shifted. Not through war, not through revolution, not through a democratic election. Through code, capital, and computing power — and through the fact that all three are concentrated in very few hands.

This book is about that shift. About what it means for the world — for states, for democracies, for the way we work, think, and trust one another. And about what it does to us: to our self-image, our sense of meaning, our understanding of what it means to be human.

But let us begin with the shift itself. And with a simple question: where exactly does the power lie?

The New Coordinates of Power

In 2024, companies worldwide invested 252 billion US dollars in artificial intelligence. That is 26 percent more than the previous year. Private AI investment in the United States alone amounted to 109 billion dollars — twelve times as much as in China and twenty-four times as much as in the United Kingdom.

[Source: Stanford HAI, AI Index Report 2025]

These numbers sound abstract. But they tell a concrete story. They tell us where the infrastructure of the future is being built and where it is not. Who owns the tools that will shape the coming decades — and who will merely be a user of those tools.

To understand this, one must look at three layers: the chips, the models, and the data. Each layer has its own power structure. And each is more concentrated than most people suspect.

Let us start with the chips. Without semiconductors, there is no AI. Every model that today generates text, creates images, or prepares decisions runs on chips manufactured in a semiconductor fabrication plant. And here the first, perhaps most astonishing concentration reveals itself: a single company — the Taiwan Semiconductor Manufacturing Company, TSMC — produces approximately 70 percent of all semiconductors manufactured on contract worldwide. For advanced chips — the high-performance processors that power AI models — the share is 90 percent.

[Source: TrendForce Q4 2025; Counterpoint Research 2025; Taipei Times, March 2026]

Seventy percent. That means: seven out of ten chips installed anywhere in the world in a server, a phone, or a car come from the factories of a single company on an island in the western Pacific. An island that is simultaneously one of the most dangerous geopolitical flashpoints on Earth.

This is not a market statistic. This is a vulnerability.

The Speed

What distinguishes this shift from earlier technological upheavals is, above all, one thing: speed.

The steam engine needed seven decades to travel from James Watt's workshop to the factories of Manchester. Electricity required roughly forty years to progress from the first light bulb to widespread infrastructure. The internet spread faster — but even there, a good twenty years passed before it fundamentally changed the world of work and the public sphere.

AI compresses this process into months.

In 2022, the smallest language model that scored above 60 percent on the MMLU benchmark — a standardized test for language comprehension — could only do so with 540 billion parameters. Two years later, in 2024, Microsoft's Phi-3-mini achieved the same score with 3.8 billion parameters. A reduction by a factor of 142. And the cost of a single query to a model at GPT-3.5 level fell in the same period from 20 dollars per million tokens to 7 cents — a decline by a factor of 280.

[Source: Stanford HAI, AI Index Report 2025, Chapter: Technical Performance]

These numbers do not merely describe technical progress. They describe a dynamic without precedent in the history of technology. A technology that simultaneously becomes more powerful, smaller, and cheaper — not over decades, but over quarters.

The consequences of this extend far beyond technology. Because speed is not neutral. It determines who can adapt and who cannot. Education systems that need decades for reforms cannot keep pace with a technology that doubles every six months. Labour markets designed for stable career biographies are overrun by a force that produces new skills faster than humans can learn them. Governments that need years to pass laws regulate a world that, by the time of passage, no longer exists.

Speed, in this context, is a form of power. Whoever is faster defines the rules of the game. Whoever is slower lives by rules that others have written.

The Invisible Threshold

There is a moment in recent history that illustrates this shift in a quiet, almost uncanny way. In November 2024 — exactly two years after the launch of ChatGPT — the volume of AI-generated content on the internet surpassed the volume of human-made content for the first time. By May 2025, the share of synthetic content had reached 52 percent.

[Source: EU Parliamentary Report EPRS 2025; referenced in Stanford AI Index 2025]

This moment went almost unnoticed. There were no headlines, no debates, no political alarm. And perhaps that is precisely the most unsettling quality of this revolution: it is invisible. It unfolds not in factories, not on battlefields, not in parliaments — but in the data centres and on the screens before which we sit every day.

We feel it. In the way search results change. In the way texts we read sound, without knowing whether a human wrote them. In the way our colleagues suddenly work faster and we wonder whether we can keep up. In the vague sense that something fundamental is shifting — without being able to name exactly what it is.

But we do not see it. And that distinguishes it from earlier upheavals. The steam engine was loud. Electricity illuminated cities. The internet connected screens. AI changes the structure of information itself — and thereby the foundation on which we understand the world.

Power as Structure, Not Intention

Two narratives dominate the public debates about artificial intelligence. One warns of a demonic technology spiralling out of control — Terminator scenarios, superintelligence, existential risks. The

other celebrates a miracle machine that solves every problem — curing diseases, eliminating poverty, driving productivity to unimaginable heights.

Both narratives are convenient because they are simple. And both miss the essential point.

The most important question about AI is not a technical one. It is not a question about whether machines can 'think' or when they will 'surpass' us. The most important question is political: who controls these systems? Whom do they benefit? And who decides the rules?

Hannah Arendt, in her 1970 essay 'On Violence,' distinguished between two forms of influence. Violence, she wrote, is instrumental — it requires tools and destroys. Power, by contrast, arises when people act together. It is never the property of an individual; it belongs to a group and exists only so long as the group holds together.

[Hannah Arendt, On Violence, 1970]

What we are witnessing today fits none of Arendt's scenarios — and makes her framework productive precisely because of that. AI is not violence in the classical sense. It does not destroy with weapons. But neither is it power in the Arendtian sense, because it does not arise from collective action. It arises from infrastructure, from ownership, from control over the tools with which reality is produced.

One might call it a third category: structural power. The power that becomes visible not in actions but in the architecture of the systems that enable or prevent actions. Michel Foucault described something similar when he wrote about the panopticon — a prison in which inmates constantly feel observed, even when no one is watching. Power operates not through the gaze but through the architecture. AI is this architecture. It decides what content we see, which job applications get through, which credit requests are approved, which medical diagnoses are suggested. It does this not because anyone means ill. It does it because the structure is built that way.

[Michel Foucault, Discipline and Punish, 1975]

What Is Different This Time — and What Is Not

There is an understandable temptation to regard every new technology as 'unique.' The history of technology is full of moments when people believed everything was changing — and when much of it then turned out to be repetition.

When the printing press reached Europe in the 15th century, the Church's monopoly on the dissemination of knowledge collapsed within a few decades. The Reformation, the scientific revolution, the emergence of an educated public — none of this would have been possible without the printing press. And yet, the liberation of knowledge was followed by religious wars, propaganda, and the manipulation of the masses.

When nuclear physics in the mid-20th century produced a new energy source and a new weapon, the geopolitical order of the world shifted within a few years. The bomb created the bipolar world of the Cold War, the balance of terror, the logic of deterrence. But it also created a form of international cooperation — disarmament treaties, arms control, diplomatic channels — that would never have arisen without the shared horror.

What is different about AI? What is repeated, and what is new?

Three things are structurally different.

First, ubiquity. Nuclear weapons were possessed by a few states. The internet reached billions, but as a communication medium. AI permeates everything simultaneously — work, health, education, military, entertainment, politics, personal relationships. There is no domain of life left untouched.

Second, invisibility. Earlier technologies were visible. You could see a factory, hear a bomb, lay a cable. AI operates in layers that do not exist for the naked eye. The algorithms that sort our news feeds, filter our job applications, or evaluate our health data are invisible to those affected. The power is there, but it does not show itself.

Third, concentration. The internet was celebrated in its early days as a decentralising force — anyone could build a website, anyone could publish. AI reverses this logic. The development of powerful models requires computing power, data, and capital on a scale that only a handful of companies can muster. The power that the internet distributed, AI gathers back.

These three qualities — ubiquitous, invisible, concentrated — do not make AI the apocalypse. But they make it a force that must be treated differently from everything that preceded it. A force that cannot be understood in the categories of the 20th century — neither as pure threat nor as pure promise.

The Question Behind the Question

This book is not a technology book. It is not an introduction to artificial intelligence, not a handbook for the digital world, not an oracle of the future.

It is a book about power and about being human. About the question of what happens when a technology that permeates everything is controlled by very few. And about the question that follows from this for each of us: what do we do with it?

The MindShift series has spanned an arc across five books — from the consciousness that creates the world, through the existential search for belonging, through the social effects of algorithmic systems, the question of professional identity, to the architecture of our perception. This sixth book

takes the greatest altitude. It looks at the world as a whole — not as an abstract system, but as the place where we all stand right now. In the spring of 2026, in the midst of an epochal turning point whose outcome is open.

In the chapters that follow, we will examine the power shifts currently underway: the contest between America, China, and Europe for AI supremacy. The new rulers who need no voters, no constitution, and no mandate. The economy of abundance that no longer distributes wages. We will see how AI changes the nature of war, how it synthesises reality, how it puts democracy under pressure.

And we will ask what can be shaped. What political tools exist. What forms of ownership are possible. And what all of this does to us — to our self-image, our sense of meaning, our capacity to live in a world that changes faster than we can understand it.

This is not a dystopia. It is a description.

And every description begins with the decision to look closely.

2

The World Game

America, China, and the Void of Europe

On 20 January 2025, the day of Donald Trump's second inauguration, a Chinese company called DeepSeek released a language model. DeepSeek-R1 was an open-source model that matched leading Western systems across numerous benchmarks — at significantly lower development costs.

The stock markets reacted. Nvidia — the company that manufactures the chips on which nearly all Western AI models run — lost 600 billion dollars in market capitalisation in a single day. It was the largest single-day loss of value in the history of stock markets.

[Source: CSIS, 'DeepSeek, Huawei, Export Controls,' March 2025; Foreign Policy, February 2025]

The message was unmistakable: the race for artificial intelligence is not an American monopoly. It is a world game. And the rules of this game are being set right now.

To understand what is happening, we must examine three models — three answers to the same question: how should power over the most powerful technology in history be organised?

The American Model: Dominance Through Freedom

The United States follows a clear line: deregulation, speed, supremacy. On 23 January 2025, the third day after inauguration, Trump signed Executive Order 14179, revoking all AI safety initiatives enacted under Biden. The rationale: AI safety is a barrier to innovation.

[Source: Executive Order 14179, 23 January 2025, Federal Register]

In July 2025, the AI Action Plan followed — 90 policy positions across 25 pages, built on three pillars: accelerate innovation, remove regulatory barriers, and secure energy supply. Not one of the 90 positions addressed questions of safety, bias, or labour market impact.

[Source: White House AI Action Plan, 23 July 2025; Sidley Austin LLP Analysis]

In December 2025, Trump went further. A new executive order targeted state-level AI legislation directly, creating a federal regulation-free zone in which companies could operate without constraint.

[Source: Executive Order, 11 December 2025, whitehouse.gov; Paul Hastings LLP Analysis]

All of this unfolded against the backdrop of a massive investment wave. In early 2025, Trump announced the Stargate project — a joint venture of OpenAI, SoftBank, and Oracle worth up to 500 billion dollars. It is intended to create AI infrastructure on a national scale.

[Source: Atlantic Council, January 2026; Stanford HAI AI Index 2025]

The American model can be summarised in a single sentence: the state clears the obstacles, capital builds. This is the embodiment of what Isaiah Berlin described in his 1958 essay 'Two Concepts of Liberty' as negative freedom: the right to be left alone.

[Isaiah Berlin, Two Concepts of Liberty, 1958]

But it has a blind spot. Negative freedom does not ask whom the fruits of that freedom benefit. It does not ask what happens when the freedom of the few becomes power over the many.

The Chinese Model: Control Through Strategy

China follows a different path. Here the state directs, and the market follows — within the boundaries the state sets. AI is not a domain of free enterprise. It is a matter of national security.

The foundations were laid early. As early as 2017, the State Council formulated a national AI strategy with a clear objective: by 2030, China would be the world leader in AI.

[Source: China State Council, New Generation AI Development Plan, 2017; CSIS Analysis, March 2025]

DeepSeek was no accident. It was the product of a strategy that treats technological autarky as a national security objective. When American export restrictions cut off access to the latest Nvidia chips, Chinese developers did not stop — they optimised.

[Source: CSIS, 'DeepSeek, Huawei, Export Controls,' March 2025]

DeepSeek's CEO Liang Wenfeng captured the dilemma in mid-2024: 'Money was never our problem. Our problem is chips.' In that one sentence, he summarised the logic of strategic circumvention: if you cannot obtain the best hardware, you must create the best software.

At the same time, China's AI ecosystem is not an open space. China's AI regulations require that every AI output align with the Communist Party line. Models that produce unwanted answers are blocked before release.

[Source: Bruegel, 'Geopolitics of AI after DeepSeek,' 2025; The Conversation, February 2025]

This is Berlin's positive freedom: freedom through collective steering. The state defines the goal — technological sovereignty, social stability, national power — and the individual serves that goal.

The strength of this model shows itself in speed. Eighty-seven percent of Chinese companies planned to increase AI spending in 2025 — compared with 78 percent of companies in the United States.

[Source: Cultura Analysis, January 2026; Atlantic Council, January 2026]

The weakness shows itself in what may not be said. An AI model that responds to a question about the Tiananmen Square massacre with 'Let's change the subject' is not merely political censorship. It is an epistemological constraint. The machine can only know what it is permitted to know.

The European Void: Values Without Machines

And Europe?

Europe regulates. In March 2024, the European Parliament passed the AI Act — the world's first comprehensive AI law. It classifies AI systems by risk level, prohibits certain practices, and establishes transparency requirements.

[Source: EUR-Lex, AI Act full text; CFR, January 2026]

This is impressive. And it is what Europe does best: writing rules, setting standards, creating ethical frameworks. The AI Act is the most ambitious regulatory project in the history of technology.

But Europe barely builds models. In global API traffic — the data flow that constitutes the actual infrastructure of the AI economy — the share of European providers is negligible.

[Source: Cultura Analysis, January 2026; Atlantic Council, January 2026]

Europe faces a dilemma that runs deeper than the question of competitiveness. It is a sovereignty dilemma. Whoever regulates a technology they do not produce depends on the goodwill of those who do.

The American ecosystem offers performance, but under a government that actively combats European regulation. The Chinese ecosystem offers a technical alternative, but one that answers to the Communist Party.

[Source: Bruegel, 2025; Garante (Italy's Data Protection Authority), January 2025]

Europe's image of the human being is that of the rights-bearer. The person is neither entrepreneur (as in the American model) nor servant of the state (as in the Chinese model), but a subject whose dignity precedes all technology. The AI Act is the legal expression of this conviction. But conviction without infrastructure is a declaration.

The Players on the Margins — and the Blocs That Are Forming

The world game has more than three players. But the rules are set by these three. Everyone else positions themselves relative to them.

India is investing heavily in the role of AI talent supplier and future producer. The IndiaAI Mission, the AI Startup Hub, and a growing digital infrastructure signal an ambition to be not merely a consumer but a shaping force.

[Source: ORF India, February 2026; Brookings, January 2026; Atlantic Council, January 2026]

Russia pursues an autarkic path under sanctions. The updated National AI Strategy through 2030 relies on domestic chips, sovereign language models, and military applications.

And then there is the alliance growing in the background. CNN reported in late 2025 on an image that reflects the new geopolitical reality: CRINK — China, Russia, Iran, North Korea. Not a formal alliance, but a network of mutual dependencies built on a common adversary: the liberal order.

[Source: CNN, December 2025; The Diplomat, December 2025; Cultura Analysis, January 2026]

Three Images of the Human, Three Futures

What collides here is more than geopolitics. It is images of the human being.

The American model sees the person as an individual who thrives best when left alone. The role of technology is to provide maximum freedom of action. The role of the state is to stay out of the way.

The Chinese model sees the person as part of a greater whole. The individual has rights, but they are subordinate to collective goals. Technology is an instrument of state power.

The European model sees the person as the bearer of inalienable rights. AI must respect these rights, regardless of economic or strategic interests.

None of these images is wrong. None is complete. The American one ignores the question of power. The Chinese one ignores the question of individual freedom. The European one neglects the question of how to protect values without possessing the means.

And all three ignore something that may be the most important of all: the question of what this technology does to the people who use it — to their self-perception, their capacity for judgement, their sense of meaning.

The world game will not be decided only in Washington, Beijing, and Brussels. It will be decided in every person who picks up their phone in the morning and wonders whether they are still living their own life — or already the life that an algorithm has deemed optimal for them.

These questions lead us to the next chapters — to the people who own the machines, to the economy they create, and to the question of what AI does to the thing we call reality.

3

The New Rulers

Tech Oligarchs and Power Without Mandate

On 20 January 2025, at the second inauguration of Donald Trump, a camera captured an image that entered collective memory: in the front row, before members of Congress, generals, and diplomats, sat four people. Elon Musk, Jeff Bezos, Mark Zuckerberg, and Sundar Pichai.

This was not a seating accident. It was a declaration. The people who control the infrastructure of the digital world — the clouds, the chips, the social networks, the AI models — were sitting where the representatives of the people normally sit. The message was clear: this power no longer hides. It shows itself.

To understand what this scene means, one must look behind the curtain. And one figure embodies the new type of power more than any other: Elon Musk.

The Musk Case — Anatomy of a New Type of Power

Musk is not merely an entrepreneur. He owns or controls Tesla (electric vehicles), SpaceX (space), Starlink (global communications infrastructure), X/Twitter (social platform), Neuralink (brain interfaces), The Boring Company (infrastructure), xAI (artificial intelligence). No individual since Rockefeller has controlled so many strategic sectors simultaneously.

In February 2025, Musk became the head of DOGE — the Department of Government Efficiency — an advisory body at the White House. Within weeks, his team reduced the federal workforce by 9 percent, froze contracts worth 38 billion dollars, and gained access to the Treasury Department's payment systems.

[Source: Military Times, 'DOGE Federal Buyouts,' March 2025; New York Times, February 2025]

What happened here goes beyond political influence. It was the direct integration of economic power into the state apparatus — without election, without mandate, without separation between business and government.

Max Weber defined the state as the institution possessing the 'monopoly on the legitimate use of physical force.' Musk does not challenge this monopoly. He bypasses it. Who needs a monopoly on violence when you control the monopoly on infrastructure?

[Max Weber, 'Politics as a Vocation,' 1919]

The Historical Parallel — and Its Limits

History has seen figures of comparable scale. The 'Robber Barons' of the Gilded Age — Rockefeller, Carnegie, Morgan — controlled oil, steel, and finance in an era when the state was weak and rules were few.

But the historical template is older and deeper. The East India Companies of the 17th and 18th centuries were not merely trading enterprises. They were private states: they minted their own currency, maintained armies, signed treaties, governed colonies. At the peak of their power, the Dutch East India Company was arguably the most powerful organisation in the world.

The parallel to the present is striking. Technology companies do not mint currency — but they control the infrastructure through which money flows. They do not maintain armies — but they supply satellite internet to combat zones. They do not govern colonies — but they determine what information billions of people receive.

Yanis Varoufakis called this phenomenon 'technofeudalism': a system in which a few platforms own the digital territory on which everyone else must work — like peasants tilling the lord's land.

[Yanis Varoufakis, Technofeudalism, 2023]

Code Is Law — and That Is a Problem

Lawrence Lessig articulated this thought as early as 1999: 'Code is law.' The architecture of digital systems determines what is possible and what is not — more precisely and more inexorably than any statute book.

[Lawrence Lessig, Code and Other Laws of Cyberspace, 1999]

An algorithm that places certain content above other content is not a technical decision. It is a political act. A platform that determines what counts as 'hate speech' and what counts as 'free speech' exercises judicial power — without a court, without the right of appeal, without a mandate.

The decisive difference from state power: state power — at least in democracies — is legitimised, controllable, and replaceable. The power of tech oligarchs is described by none of these qualities. It is legitimised by the market, controllable only conditionally, and replaceable only in theory — because the infrastructures they control have become too fundamental to abandon.

This is not a conspiracy. It is a structure. And that is precisely why it is so difficult to do anything about it.

What Follows from This

The rise of the tech oligarchs is not an argument against technology. It is an argument for a political architecture commensurate with the actual power.

The antitrust legislation of the 19th century was a response to the robber barons. The social legislation of the 20th century was a response to industrialisation. The question of the 21st century: what institutions do we need for a world in which private individuals control the infrastructure of knowledge, communication, and decision-making?

This question is not rhetorical. It is urgent. And it leads us to the next chapter — to the economy that this technology creates, and to the question of who still earns anything in it.

4

The Economy of Abundance

and Why It Doesn't Work

In 1930, John Maynard Keynes wrote an essay entitled 'Economic Possibilities for Our Grandchildren.' He predicted that a hundred years hence — around 2030 — productivity would have grown to such an extent that people would need to work only fifteen hours a week. Technology, Keynes wrote, would solve the problem of scarcity. The remaining time would become time for culture, contemplation, and human development.

[John Maynard Keynes, 'Economic Possibilities for Our Grandchildren,' 1930]

We are approaching the date Keynes named. Productivity has grown as he predicted — and has significantly exceeded his forecast. But the fifteen-hour workweek has not materialised. On the contrary: most people work more than ever, and many are not certain whether they will have a job at all in five years.

What went wrong?

The Paradox of Abundance

AI possesses a quality without precedent in economic history: it reduces the marginal cost of cognitive labour to near zero.

The cost of inference — a single query to an AI model — falls by approximately a factor of ten each year. What cost twenty dollars per million tokens in 2022 had dropped below seven cents by 2025. This is not a gradual improvement. It is a collapse — akin to what happened with information after the emergence of the internet, but this time concerning thought itself.

[Source: Stanford HAI AI Index 2025; Epoch AI, 'Inference Cost Trends,' 2025]

When the cost of thinking approaches zero, much of what previously justified the payment of human labour becomes economically devalued. The text a journalist wrote. The analysis a consultant prepared. The code a developer created. The sketch a designer drew. The diagnosis a doctor made.

This does not mean these professions will disappear. It means the economic foundation on which they stand is eroding.

The Numbers — and What They Mean

The World Economic Forum predicts in its Future of Jobs Report 2025 that by 2030, AI and automation will create 170 million new jobs worldwide while simultaneously eliminating 92 million. Net gain: 78 million. This sounds reassuring.

[Source: World Economic Forum, 'Future of Jobs Report 2025']

But behind the aggregate numbers lies a rupture. The new jobs emerge in different fields, require different qualifications, and are often in different countries than those that disappear. The truck driver who loses his job to autonomous driving will not become an AI specialist. The accountant whose work is taken over by an algorithm will not retrain as a machine learning engineer.

A 2025 study by the Pew Research Center paints a more nuanced picture: 61 percent of Americans believe AI will have a 'serious impact' on jobs in the next twenty years — but only 28 percent believe their own job will be affected. The classic perception gap: they see the threat — but always in others.

[Source: Pew Research Center, 'AI and the Future of Work,' 2025]

The Labour Share and the Schumpeter Paradox

A deeper trend is visible in the macroeconomic data. In most advanced economies, the share of wages in GDP — the so-called labour share — has been declining for decades. AI accelerates this trend, because each productivity leap accrues primarily to the owners of capital — those who own the models, the chips, and the infrastructure.

Joseph Schumpeter described in 1942 the mechanism of 'creative destruction': new technologies destroy old industries and create new ones. This process is painful but ultimately productive — because the new is more powerful than the old.

[Joseph Schumpeter, Capitalism, Socialism and Democracy, 1942]

The paradox of our time is that the creative part may be slowing while the destructive part accelerates. When AI acquires in months what takes humans years to learn, the window of

adaptation collapses. Schumpeter assumed that people would have time to restructure. What if they do not?

The Distribution Question

The fundamental problem is not productivity. The problem is distribution.

When AI increases productivity by 40 percent — as many studies claim for certain tasks — the question arises: who gets the gain? The worker, who now works faster? The company, which now needs fewer workers? Or the AI developer, whose model makes it possible?

At present, the answer is clear: the lion's share goes to those who own the technology. The five largest technology companies in the world reached a combined market capitalisation of twelve trillion dollars in 2024. The combined wealth of the ten richest tech billionaires grew by 37 percent in a single year.

This is not a bug in the system. It is the logic of the system. A technology that reduces the marginal cost of cognitive labour to zero, under unchanged ownership conditions, leads to a concentration of wealth without precedent.

The Old Dream and the New Question

Keynes was not wrong in his productivity forecast. He was wrong in the tacit assumption that rising productivity automatically leads to the common good. He did not foresee a world in which technology is so productive that it no longer needs most people — and in which the fruits of that productivity accrue to a few.

The question that the AI economy poses is not technical. It is the oldest question of political economy: who owns what? And who decides?

This question will lead us later — in Chapter 9 — to political answers: to universal basic income, to new forms of ownership, to the question of whether a just economy can be built at all in a world where machines do a substantial share of the work.

But first we must look at the other side of this shift. At the question of what AI does to what we consider reality. To truth, to war, and to our capacity to distinguish one from the other.

PART II

The Destruction of Reality

What AI does to truth, trust, and democracy

5

War Without a Battlefield

How artificial intelligence changes the nature of war

In February 2026, for the first time in history, humanoid robots were deployed to an active combat zone. Two units of the Phantom MK-1, developed by the Dallas-based startup Foundation Labs, were delivered to the Ukrainian military. The robots are equipped with AI-based object recognition and can operate autonomously at distances of up to two kilometres.

[Source: TIME, 9 March 2026; Military Times, 13 March 2026; Foundation Labs]

At the same time, one week later, Ukraine opened its war data to allied states and defence companies: millions of annotated video recordings, damage assessments, and detection patterns — the largest open real-battlefield dataset ever assembled.

[Source: Military Times, 13 March 2026; Ukrainian Ministry of Defence]

And while the future of war is being tested in Ukraine, its present is already operating in Gaza: AI systems that issue target recommendations for air strikes, with minimal human review time and maximum speed.

This chapter is about the transformation of war. Not as an abstract threat, but as something that is happening right now — in two conflicts that will define the future of armed violence for decades to come.

Ukraine: The Laboratory

The war in Ukraine has become the world's most important testing ground for AI-assisted warfare. Not because the technology was invented there, but because it is being applied there — under the pressure of real combat, with real casualties.

The numbers are both impressive and sobering. Ukraine has more than five million drones in service. In January 2026 alone, military authorities conducted more than 4,000 unmanned aerial vehicle operations daily.

[Source: TIME, March 2026; CSIS, March 2025; Modern War Institute, West Point]

AI plays an increasingly central role. Neural networks in Ukraine's DELTA system automatically identify ground and air targets in real time. Automated targeting management in the Ukrainian Armed Forces has reduced the cycle from detection to strike from hours to minutes.

[Source: CSIS, 'Ukraine's Future Vision for AI-Enabled Autonomous Warfare,' March 2025]

The result is measurable: the integration of AI into first-person-view drones has raised the hit rate from 30 to 50 percent to approximately 80 percent. Four out of five drones reach their target.

[Source: US Army War College; West Point Lieber Institute, 2025]

But it is the next step that raises the ethical questions. AI-assisted drones in Ukraine already fire autonomously when communication with the operator is lost. This is technically understandable — in an electronically jammed environment, connections are regularly interrupted. But it means: the machine makes the decision to kill. Without a human in the loop.

[Source: TIME, March 2026; West Point Lieber Institute, February 2025]

And Europe is arming up. In March 2025, the European Commission called for an 'unprecedented wave of investment in European defence' and designated AI as a 'critical defence technology.' Anduril — founded by Palmer Luckey, the creator of Oculus VR — supplies autonomous systems and has signed contracts with multiple NATO countries.

[Source: MIT Technology Review, January 2026; European Commission, March 2025]

Gaza: The Precedent

While Ukraine shows how AI shapes the war of the future, Gaza shows what happens when AI systems make decisions about human lives in the present.

In April 2024, an investigation by the Israeli-Palestinian magazine +972 and Local Call, based on testimony from six Israeli intelligence officers, revealed a system called 'Lavender.' Lavender was an AI programme that flagged individuals in the Gaza Strip as suspected militants — based on communication patterns, social connections, and behavioural data. According to military estimates, Lavender flagged approximately 37,000 Palestinians as potential targets.

[Source: +972 Magazine, April 2024; The Guardian, 2024; Human Rights Watch, September 2024]

Alongside Lavender operated 'The Gospel' (Habsora), a system that identified buildings and structures as bombing targets. And 'Where's Daddy?' — a programme that tracked the movements

of specific individuals and alerted when they returned home, so that the residential building could be struck.

The testimony of the intelligence officers paints a picture that blurs the line between human control and machine decision-making. The average review time for a Lavender recommendation was approximately 20 seconds per target.

[Source: +972 Magazine, April 2024; AOA Analysis, November 2025]

The consequences were devastating. The Gospel generated up to 100 targets per day — where human analysts had previously identified approximately 50 per year. This leap called into question one of the central pillars of international humanitarian law: the principle of proportionality.

[Source: +972 Magazine, April 2024; RUSI Analysis, July 2024]

UN Secretary-General Guterres declared himself 'deeply troubled.' Human Rights Watch analysed the systems and concluded that they 'undermine the foundations of international humanitarian law.'

[Source: UN/Guterres, 2024; Human Rights Watch, September 2024; The Guardian/Microsoft, September 2025]

Israel disputes parts of these accounts and emphasises that the systems serve as decision support, not autonomous weapons. The debate is not settled. But the precedent has been set: AI in a real conflict, in real time, with real casualties.

The Disappearing Line

Under international law, three principles govern the use of military force: distinction between combatants and civilians, proportionality between military necessity and civilian harm, and accountability for every use of force.

[Michael Walzer, Just and Unjust Wars, 1977]

AI systems subject each of these principles to a new test.

Distinction: Can an algorithm trained on communication patterns and movement data reliably distinguish a combatant from a civilian who simply lives in the same neighbourhood?

Proportionality: If a system generates a hundred targets per day where humans identified fifty per year, and if human review time is 20 seconds — does meaningful control still exist?

Accountability: If a drone fires autonomously because the connection to the operator is lost — who bears responsibility for the damage? The programmer? The commander? The machine?

The Failed Regulation

Austria's Foreign Minister Schallenberg called the situation 'the Oppenheimer moment of our generation.' The comparison is evocative but limps in one respect: Oppenheimer worked for the state. Today's developers work for startups.

The United Nations has tried. Secretary-General Guterres demanded a legally binding treaty by 2026 prohibiting autonomous systems without 'meaningful human control.' The treaty did not materialise.

[Source: UN/ICRC, 2024; Geneva Negotiations; TRENDS Research, 2025]

The largest military powers — the US, Russia, Israel — are blocking. And the reality on the battlefield is outrunning the negotiations. Every week in which no treaty is signed, new facts are created.

What Is at Stake

In Corbera — a village on a rocky hill in Catalonia — there is a moment that frames this story in a quiet, compelling way. In March 2025, the local municipality held a discussion on the question: should the use of AI-assisted drones be permitted on municipal territory? Not military drones — drones for agriculture, forestry, emergencies. The debate lasted three hours.

[Source: MIT Technology Review, January 2026]

Ukraine is the new testing ground. What is developed, tested, and optimised there will shape the wars of the coming decades. The drones that fly in Ukraine today will fly everywhere tomorrow.

And Gaza shows what happens when AI systems are applied to human lives in a real conflict, with minimal control and maximum speed. The question is: do we want to live in a world where machines decide who lives and who dies?

The answer to this question determines not only the future of war. It determines the future of morality. For a society that accepts that algorithms make decisions about life and death has already changed its understanding of what it means to be human.

6

The Synthetic Public

Deepfakes, disinformation, and the end of shared reality

In November 2024, two years after the launch of ChatGPT, something happened that went almost unnoticed: the volume of AI-generated articles on the internet surpassed the volume of human-written articles for the first time. By May 2025, 52 percent of all web articles were of synthetic origin.

[Source: Graphite/Five Percent Research, October 2025; confirmed by Axios, eWeek, Vice, Yahoo Tech]

The moment came without fanfare. There was no press conference, no political alarm, no date anyone wrote in a calendar. And perhaps this is precisely the most disturbing quality: the threshold was crossed — and no one noticed.

Europol warned in a report that by 2026, up to 90 percent of online content could be synthetically generated. Stanford University confirmed the trend. The University of Cambridge found in 2024 that humans could no longer reliably distinguish AI texts from human ones.

[Source: Europol Innovation Lab; Stanford HAI AI Index 2025; Cambridge University, 2024]

What does it mean for a society when more than half of what it reads was written by machines? Not all of it is disinformation. Most is product descriptions, SEO texts, automated reviews. But the line between harmless generation and targeted manipulation is blurring.

The Elections That Can No Longer Be Trusted

In January 2024, voters in the US state of New Hampshire received a phone call. The voice apparently belonged to President Biden. It urged people not to vote in the primaries. The voice was AI-generated. The Federal Communications Commission fined the responsible party 6 million dollars — but the call reached thousands of people.

[Source: Brennan Center for Justice; CNN, 2024]

Since then, the situation has not calmed down. It has normalised.

In Romania, AI-generated content influenced the December 2024 presidential election so massively that a court annulled the results. Across Europe, AI systems were detected in more than ten major election campaigns since early 2024 — from Slovakia to Indonesia, from Bangladesh to South Africa.

[Source: CETaS/Alan Turing Institute, November 2025; CIGI, 2025; EU Parliamentary Report EPRS 2025]

And in March 2026, the National Republican Senatorial Committee in the United States published a deepfake video of Democratic Senate candidate James Talarico from Texas containing synthetically generated, false statements.

[Source: CNN, 13 March 2026]

The most remarkable thing about it was not the technology. It was the reaction — or rather, its absence. The Republican Senate organisation saw no problem. The Democratic Party protested, but there were no legal consequences.

The Liar's Dividend

There is a twist in this story that may be more dangerous than the deepfakes themselves. Researchers call it the 'Liar's Dividend' — the benefit that liars derive from the very existence of deepfakes. If everything can be fake, then what is real can also be declared fake.

[Source: Chesney & Citron, 'Deep Fakes: A Looming Challenge'; Brennan Center for Justice]

The Liar's Dividend does not destroy truth through forgery. It destroys the very possibility of distinguishing truth from forgery. A politician caught in corruption can declare the evidence a deepfake. A government filmed using violence can claim the video is fabricated.

According to a survey by IE University, 31 percent of Europeans believe that AI influenced their voting decision. And 40 percent are concerned that they can no longer distinguish real content from generated content.

[Source: IE University, European Tech Insights, October 2024]

The Poisoned Sources

Deepfakes are the visible tip. Beneath them lies a broader, subtler transformation: the poisoning of the information sources themselves.

Before the Australian parliamentary elections in May 2025, investigators discovered a pro-Russian campaign that systematically fed AI chatbots false information — so that they would present it as fact when answering voter questions.

[Source: ABC News Australia, May 2025; CETaS/Alan Turing Institute, November 2025]

In the United Kingdom in 2024, roughly one in eight voters turned to AI chatbots for election information. In several cases, the bots provided incorrect information — about voting dates, constituencies, and candidate positions.

[Source: CETaS/Alan Turing Institute, November 2025]

The fifth book in the MindShift series, *The Quiet Shift*, examined the architecture of perception: how feeds, algorithms, and personalised recommendations create an individual reality that everyone takes to be objective. The synthetic public goes one step further: it replaces not merely the presentation of information, but the information itself.

Arendt, Factual Truth, and What We Are Losing

Hannah Arendt wrote an essay in 1967 titled 'Truth and Politics.' In it she distinguished between two kinds of truth: rational truth — mathematical and logical propositions that cannot be refuted — and factual truth — historical and empirical claims that depend on witnesses, documents, and the willingness to preserve them.

Factual truth, Arendt wrote, is more vulnerable than rational truth. For it depends on people who remember it — and people forget, are bought, or are intimidated.

[Hannah Arendt, 'Truth and Politics,' 1967]

This is precisely what is happening now. The synthetic public does not destroy individual facts. It destroys the field on which facts are recognised as such. When every video can be fake, when every text can be machine-made, when every phone call can be synthesised — it is not just a particular lie that disappears. The ground on which truth stands disappears.

Arendt warned: where there is no factual truth, there is no space for politics. For politics presupposes the existence of a shared world about which one can argue. If this shared world no longer exists — if everyone lives in their own algorithmically curated reality — then the subject of argument vanishes too.

The Legislation That Cannot Keep Up

The political responses are fragmented and slow. In the United States, 26 states have passed laws on deepfakes in elections — but most are limited to labelling and do not provide effective sanctions.

[Source: Ballotpedia; Paul Hastings LLP; Cornell Law School, 2025]

In the EU, the AI Act is to require from August 2026 that AI-generated content be labelled. But implementation is delayed: the European Commission is behind schedule, industry is requesting extensions.

[Source: EU Parliamentary Report EPRS 2025; European Commission, November 2025; CFR, January 2026]

Technology moves faster than any law. And the more fundamental question is whether laws are even the right instrument. Labelling deepfakes will not resolve the crisis of trust if trust itself has already been destroyed.

What Remains

Guy Debord wrote in 1967 — the same year as Arendt — *The Society of the Spectacle*. He described a world in which the image replaces reality.

[Guy Debord, The Society of the Spectacle, 1967]

Debord could not have imagined how right he would be. In his world, the spectacle was created by advertising, television, and consumption. In ours, it is created by algorithms that generate content individually tailored for each person — the perfect spectacle, which every viewer takes for reality.

What remains when shared reality vanishes?

What remains is the question: can democracy — that form of coexistence premised on the assumption that citizens can make informed decisions — survive in a world where information itself has become unreliable? This question leads us to the next chapter.

Surveillance as a Business Model

The transparent human between comfort and control

There is a scene that repeats millions of times each day, all over the world. A person opens their phone. Unlocks it with their face. Scrolls through a news feed, checks a navigation route, sends a voice message, makes a contactless payment, searches a symptom on Google.

With each of these actions, a data point is generated. Together, they form a profile — more finely grained, more comprehensive, and more current than anything a secret service of the past century could have collected. And unlike the work of a secret service, it is created voluntarily.

This chapter is about the merger of service and surveillance. And about the possibility that the distinction between the two was never as clear as we would like.

The Explicit System: China

Let us begin with what is most visible — and most misunderstood.

China's Social Credit System is often portrayed in the West as an Orwellian nightmare: a single score for every citizen that determines airline tickets, schools, and housing. The reality is more complex — and no less troubling upon closer inspection.

There is no single nationwide citizen score. The original planning sketch of 2014 did not mention scores at all. What exists is a patchwork of regional and sectoral systems: credit ratings for businesses, administrative lists for individuals, local pilot projects with varying criteria.

[Source: China Law Translate, April 2025; Newsweek, November 2025; MSA Advisory, 2026; ChoZan, November 2025]

For individuals, the system operates primarily through blacklists. By September 2025, approximately 200,000 additional people had been placed on such lists — not for political speech, but for failure to comply with court orders, non-repayment of debts, violation of professional obligations.

[Source: Wikipedia/Social Credit, September 2025; China Law Translate, April 2025]

The March 2025 directive — 23 measures issued by the Central Committee and the State Council — reinforced this course: standardisation of criteria, expansion of data sharing between agencies, new sanction mechanisms.

And this is the decisive point. What is troubling about China's system is not what it does today. It is what it could do tomorrow. The infrastructure is in place. The cameras are installed. The data flows. The path from credit rating to political control is not a long one — it requires only a political decision.

The Implicit System: The West

And now the part that is closer to us — and that we are less eager to examine.

In Western democracies, there is no 'Social Credit System.' But there are credit scores that determine whether someone gets a loan, an apartment, or a mobile phone contract. There are algorithms that decide which job postings are shown to us and which are not. There are health platforms that collect health data and sell them to insurance companies.

None of these systems calls itself surveillance. They call themselves service, efficiency, personalisation, risk management. But in sum, they create a structure functionally close to what China is building deliberately.

Shoshana Zuboff coined the term 'surveillance capitalism' for this in 2018: an economic form in which human experience is extracted as free raw material and processed into behavioural predictions that are sold on the market.

[Shoshana Zuboff, The Age of Surveillance Capitalism, 2018]

Since 2018, nothing has changed about this fundamental logic. What has changed is the capability of the systems. AI models make predictions more precise. Facial recognition becomes cheaper. Voice assistants hear not only commands but moods.

The difference from China is not the absence of surveillance. The difference is its form. China surveils explicitly, through the state, visibly. The West surveils implicitly, commercially, comfortably.

What Observation Does to the Self

The first book in the MindShift series, 'You Are the Universe,' posed a question that returns here: how does reality come into being? The answer was: nothing in our perception is given directly. We construct the world — from sensory impressions, experiences, expectations, the stories we tell ourselves.

What happens when this construction is constantly observed, measured, and evaluated from outside?

Byung-Chul Han described in *Psychopolitics* the shift from the disciplinary society to the control society: a society in which power operates not through prohibitions but through seduction. Not 'you must' but 'you can.' The system does not compel — it optimises. And what it optimises is us.

[Byung-Chul Han, Psychopolitics, 2014]

Han called it 'voluntary self-exposure.' The person shares their data not under duress but from a desire for comfort, recognition, belonging. The smartphone is not a handcuff. It is a mirror that shows us an optimised version of ourselves.

Jeremy Bentham's panopticon, which Foucault turned into a metaphor for modern power, operated through the fear of being observed. The new panopticon operates through the pleasure of being seen. Instagram, TikTok, LinkedIn — these are not prisons. They are stages. But the effect is analogous: we behave as the system expects.

The European Response — and Its Limits

Europe has attempted a third way. The GDPR, in force since 2018, gave individuals rights over their data: access, deletion, objection, portability. The AI Act supplements data protection with transparency requirements.

This is more protection than any other region in the world offers. But it is also less than it appears at first glance. The GDPR did not stop data collection — it regulated it. The business model remained the same: data is the oil, the user is the source.

And Europe also surveils. Facial recognition is being tested in police departments. Predictive policing software calculates in several countries where crime is most likely to occur. The difference from China is one of degree, not of principle.

The Question Everyone Must Ask Themselves

The Stasi — the secret police of East Germany — at the height of its power employed approximately 91,000 full-time and some 189,000 unofficial collaborators. They surveilled 16 million citizens. The ratio was roughly 1 to 57.

Today, a single smartphone generates more data in a day than a Stasi officer could have collected in a month. And it does so voluntarily, without pressure, without a conscious decision.

Edward Snowden revealed in 2013 what state intelligence agencies do with this capacity. Since then, the field has shifted: the state is no longer the primary collector. The primary collector is the market.

Foucault's panopticon was a building. A prison with a tower in the centre from which every inmate could be observed at any time. The point was not actual observation — it was the possibility of observation. The inmate behaved as if watched — because they could not know for certain.

[Michel Foucault, Discipline and Punish, 1975]

The question everyone must ask themselves is not: am I being surveilled? That question has been answered. The answer is: yes. For years.

The real question is: what does it do to me? What changes in my behaviour, my thinking, my self-image, when I know — or suspect — that every action of mine is recorded, analysed, and evaluated?

The answer is: it makes us more cautious. More conformist. More predictable. Not because anyone demands it. But because the system is built so that adaptation is rewarded and deviation is penalised.

Between the Chinese system that surveils explicitly and the Western system that surveils implicitly, there is no abyss. There is a spectrum. And the question of where on this spectrum we want to live is one of the most important questions of our time.

The democracy at the centre of the next chapter requires informed, autonomous citizens. Citizens who think freely and decide freely. But can transparency — the knowledge or suspicion of being watched — coexist with freedom?

That is the question Chapter 8 asks.

8

Democracy in Intensive Care

Why the old forms cannot withstand the new pressure — and what might come next

On 17 March 2026 — three days before the time at which these lines are being written — the V-Dem Institute at the University of Gothenburg published its annual democracy report. It analyses the state of democracy in 180 countries on the basis of more than 500 indicators collected by thousands of experts worldwide.

The finding: the United States was classified for the first time in over fifty years not as a 'liberal democracy' but merely as an 'electoral democracy' — a category indicating that elections are held but fundamental freedoms, the rule of law, and the separation of powers are significantly constrained.

[Source: V-Dem Institute, Democracy Report 2026, 17 March 2026; CNN, 18 March 2026]

A few days earlier, Freedom House had published its own annual report: the twentieth consecutive year in which global freedom declined. Fifty-two countries recorded deterioration, only 22 recorded improvement.

[Source: Freedom House, Freedom in the World 2026, March 2026]

And a survey by the Pew Research Center from spring 2025 showed: 62 percent of Americans were dissatisfied with how democracy functioned in their country. Among young adults (18–29), the figure reached 72 percent.

[Source: Pew Research Center, Spring 2025; published February 2026]

This chapter could have been written in any earlier decade. The erosion of democracy is not a new topic. But what makes the current moment different is the specific character of the threat. It is not tanks before parliament. It is not a general seizing power. It is the slow, invisible destruction of the conditions that make democracy possible.

The Three Preconditions

Democracy rests on three preconditions so self-evident that they are rarely spoken aloud.

First: informed citizens. Democracy presupposes that people who participate in elections have a reasonably reliable basis on which to form judgements. They need not be experts — but they must have access to information that allows them to distinguish reality from propaganda.

Second: shared facts. Democracy does not require that everyone agrees. But it requires that a shared reality exists about which disagreements can be conducted. If there are no facts accepted by all sides, then debates become ritual.

Third: public discourse. Democracy lives through debate — through the possibility of testing arguments, comparing positions, finding compromise. This presupposes a space in which different voices can be heard — not only the loudest.

AI undermines all three preconditions simultaneously. That is what is qualitatively new.

The synthetic public described in the previous chapter attacks the first precondition: informed citizens. When half the content is machine-generated and circulated without quality control, the foundation for informed decisions collapses.

The Liar's Dividend, also described in the previous chapter, attacks the second precondition: shared facts. If everything can be faked, then what is real can also be declared fake. There is no common factual ground left.

And the algorithmic curation analysed in *The Quiet Shift* attacks the third precondition: public discourse. When every person receives their own algorithmically optimised feed, the shared space for discussion no longer exists.

The Midterms as Test Case

The US midterm elections in November 2026 will be the first major vote in which all three dynamics act simultaneously on a democracy already weakened by institutional pressure.

Deepfakes are already in use — not by anonymous actors, but by one of the two major party organisations, the National Republican Senatorial Committee.

[Source: CNN, 13 March 2026; Ballotpedia; Cornell Law School]

At the same time, the Trump administration has systematically undermined state-level AI regulation through its executive orders and steered development into a regulation-free zone.

And the platforms on which the campaign is conducted are controlled by the actors we described in Chapter 3: owners without mandate, without accountability, without democratic legitimacy.

Habermas and the Destroyed Public Sphere

Jürgen Habermas described in 1962, in *The Structural Transformation of the Public Sphere*, how in the 18th and 19th centuries a public space emerged — coffeehouses, salons, newspapers — in which citizens exchanged arguments as equals. This 'bourgeois public sphere' was the precondition of liberal democracy: a place where opinions were shaped not by power but by reason.

[Jürgen Habermas, The Structural Transformation of the Public Sphere, 1962]

Habermas diagnosed its decline even then: through the commercialisation of media, through manipulation by advertising and public relations, through the retreat of citizens into the private sphere.

The algorithmic public sphere is not a public sphere in the Habermasian sense. It is an attention marketplace in which not arguments but emotions prevail. Not maturity but engagement. Not truth but clickability.

What Radio and Television Could Not Do

It would be comforting to view the current crisis as repetition — yet another wave of technological change that democracies have weathered before.

And it is true: new media have challenged democracies before. Radio made possible both Hitler's mass manipulation and Roosevelt's Fireside Chats. Television changed politics — the 1960 Kennedy-Nixon debates demonstrated that appearance could defeat substance.

The difference is threefold.

First, speed. Radio took decades, television took decades. AI-generated disinformation takes seconds. The speed cycle to which society must adapt has surpassed the processing speed of human institutions.

Second, personalisation. Radio and television broadcast the same thing to everyone. AI-powered platforms broadcast something different to everyone. Manipulation has become individual.

Third, uncontrollability. Radio could be switched off, newspapers banned, broadcast licences revoked. Algorithmic feeds are global, decentralised, and adaptive.

Is This the End?

It would be wrong to write this chapter as a eulogy. Democracies have emerged from worse crises. The Weimar Republic collapsed — but the Federal Republic arose from its ashes. The post-war consensus fell apart — but the European Union created a new framework. Each time the form changed. The values endured.

The question, therefore, is not: is democracy finished? The question is: are the existing forms sufficient?

Democracy as we know it was designed for a world in which information flowed slowly, in which the public sphere was structured by a few media outlets, and in which power was tied to territory. None of this is true any longer.

This does not mean that democratic values — equality, freedom, accountability, human dignity — are obsolete. They are more urgent than ever. But the institutions that are meant to embody them need renewal.

John Dewey wrote in 1927, in *The Public and Its Problems*: democracy is not an institution. It is a living practice — the way people organise themselves, deliberate, and make decisions. This practice does not disappear because the media change. But it needs new forms in order to remain alive.

[John Dewey, The Public and Its Problems, 1927]

That is the task that lies before us. Not the defence of old institutions for their own sake — but the invention of new ones that embody the same values in a radically changed world.

The Turning Point

Here the diagnosis ends. The first eight chapters of this book have described what is happening: the shift of power, the world game of the great powers, the rise of the tech oligarchs, an economy that produces abundance but does not share it, war without a battlefield, the synthetic public, surveillance as a business model, and democracy under pressure.

It was important to look precisely. No embellishment, no dramatisation. The facts, as they are, in March 2026.

But this book is not a tombstone. It is a compass. And a compass does not only show where you are. It shows where you can go.

The final four chapters are about that: the return of the political, the architecture of freedom, the question of what it means to be human when machines can do more and more — and a letter to a world in transition.

The diagnosis is made. Now the shaping begins.

PART III

The Shaping

What is possible — and what depends on us

9

The Return of the Political

Why technology alone solves nothing

There is an illusion so widespread that it is scarcely recognised as one: the illusion that technical problems require technical solutions. That if AI creates problems, better AI will fix them. That it is enough to write the right algorithm — and the system will work.

This illusion is convenient because it avoids the harder question. The harder question is not: how do we make AI better? It is: who decides? Who owns the value that automation creates? Who bears responsibility for the damage? Who draws the boundaries?

These are not technical questions. They are political questions. And the fact that they are treated as political questions is the precondition for the decades ahead not being determined solely by the interests of those who own the machines.

What Europe Is Attempting

The EU AI Act is the most ambitious experiment in AI regulation undertaken to date. Passed by the European Parliament in March 2024, it classifies AI systems by risk level, prohibits certain practices — such as real-time biometric surveillance in public spaces — and establishes transparency requirements.

[Source: EUR-Lex, AI Act full text; CFR, January 2026]

This is impressive. And it is insufficient.

Insufficient not because the rules are bad, but because regulation alone does not solve the fundamental problem. You can regulate the use of a technology — but if you do not own the technology itself, you depend on the goodwill of those who do.

And the implementation reveals the limits: in November 2025, the European Commission proposed postponing certain deadlines — transparency requirements for generative models were deferred, sanctions diluted.

[Source: European Commission, Digital Omnibus, November 2025; EU Parliamentary Report EPRS 2025]

And yet: the AI Act exists. It sets a framework. It demonstrates that democratic societies possess the capacity to politically steer technological development — even if the result is imperfect.

What the Pilots Show

Alongside regulation, there is a second level of political design: distribution. If the old mechanism — labour in exchange for wages — is no longer sufficient to ensure participation for all, new forms are needed. Is there evidence that they work?

The answer is: yes. Limited, provisional, but yes.

In April 2025, DIW Berlin and the Vienna University of Economics and Business published the results of the German Basic Income pilot project — the longest and methodologically most rigorous randomised study of unconditional basic income in Europe. One hundred and twenty-two people received 1,200 euros per month for three years.

[Source: DIW Wochenbericht 15/2025; UBIE/BIEN, April 2025; Mein Grundeinkommen e.V.]

The results: no withdrawal from the labour market. Participants continued to work, often more deliberately. Mental health improved significantly. Stress and anxiety declined. Subjective life satisfaction rose — and remained elevated even after payments ended.

Catalonia launched one of Europe's most ambitious pilots in 2024: 5,000 individuals receive approximately 800 euros per month (adults) and 300 euros (children), unconditionally, for two years.

[Source: Borgen Project, May 2025; Wikipedia/UBI pilots; Government of Catalonia]

Finland tested a basic income of 560 euros per month for 2,000 unemployed individuals in 2017–2018. The employment effect was modest, but the gains in well-being were significant. The Marshall Islands launched a climate basic income pilot in autumn 2025 — for climate migrants.

[Source: Wikipedia/UBI; Chiang Rai Times, December 2025; Marshall Islands Government, November 2025]

What do these pilots show? Three findings that emerge consistently across countries, cultures, and income levels:

First: People do not stop working. They work differently — more deliberately, more selectively, often in areas that matter more to them.

Second: Mental health improves. Stress declines. Security changes how people think and decide.

Third: The effects are real but limited to the conditions of the pilot. A three-year test with 107 people is not proof that a national programme would work the same way. Scaling is a separate challenge.

The Stance Behind the Instrument

It would be possible to write this chapter as a catalogue of solutions: automation tax, data dividend, commons models, public stack, unconditional basic income. All of these instruments are discussed, some are being tested, none has been implemented at scale.

But that would be a different book. And it would miss the decisive point.

The decisive point is not the instrument. It is the stance behind the instrument. For every political tool presupposes an answer to a deeper question: what conception of the human being does it embody?

An automation tax presupposes that the value created by machines also belongs to those whom machines displace. A data dividend presupposes that data is not free raw material but a contribution deserving of reward. A basic income presupposes that every human being has a right to material participation — unconditionally, simply because they exist.

None of these presuppositions is self-evident. Each requires a conscious decision — a political decision that cannot be derived from technology, only from a conception of the human being.

And that is precisely the return of the political: the recognition that the most important questions of the AI era are not questions for engineers. They are questions for all of us.

Beyond the Market, This Side of Utopia

Karl Polanyi described in 1944, in *The Great Transformation*, how society in the 19th century had to learn to re-embed the unleashed market in social structures — through laws, institutions, social systems. This process was painful, protracted, and politically contentious. But the alternative was destruction.

[Karl Polanyi, The Great Transformation, 1944]

We face a similar task. The AI economy — with its marginal costs approaching zero, its concentration dynamics, its decoupling of productivity from wages — needs new rules. Not out of hostility to technology. Out of respect for human beings.

Elinor Ostrom received the Nobel Prize in Economics in 2009 for her research on commons — collectively managed resources that need be neither privatised nor nationalised to be used efficiently. Forests, fisheries, irrigation systems — Ostrom showed that communities can manage complex resources sustainably when certain conditions are met: transparency, participation, accountability, adaptability.

*[Elinor Ostrom, *Governing the Commons*, 1990]*

Whether AI can become a commons — a collectively managed infrastructure, similar to public libraries, road networks, or the internet in its early form — is an open question. But it is a question that deserves to be asked before private ownership becomes so entrenched that alternatives appear unthinkable.

What Must Happen Now

The return of the political does not mean that politicians have all the answers. It means that the questions AI raises must be negotiated publicly — not decided in the boardrooms of a few companies.

This sounds old-fashioned. And perhaps that is precisely its strength. In a world driven by speed, by exponential curves and quarterly reports, a commitment to slow, democratic, imperfect processes is a radical act.

The question is not whether technology is faster than democracy. It obviously is. The question is whether we are prepared to accept the slowness of democracy as a price worth paying — because the alternative is not speed. The alternative is power without legitimacy.

The next chapter asks to whom the future belongs — and whether it is possible to build an architecture of freedom that can withstand the pressure.

The Architecture of Freedom

Ownership, access, and the question of who owns the future

If this book had to be distilled to a single question, it would be this: who owns artificial intelligence?

Not: how intelligent will it become? Not: when will it surpass us? But: who owns the models, the data, the infrastructure? Who decides what is done with the value these systems create? And by what rules?

The answer to this question determines in what world we will live. It is not the technology that decides. It is the ownership.

Three Architectures

Three models are taking shape that could define the coming decades. None is purely theoretical. Each already exists in embryonic form, and each embodies a particular image of the human being.

The first model: private infrastructure. A few companies own the models, control the platforms, operate the cloud. Access is for a fee, conditions are dictated by the provider. This is the model of maximum efficiency — and maximum dependency.

The second model: state control. The state owns or directs the infrastructure. It determines who gets access, under what conditions, for what purposes. This is the model of maximum control — and minimum freedom.

The third model: shared infrastructure. The fundamentals — foundation models, computing capacity, data standards — are treated as a public good. On top of them, private and state actors operate. This is the model of maximum participation — and the greatest complexity.

These are not abstract scenarios. The first model is the current reality of the Western AI economy. The second is China's approach. The third is a vision still searching for its realisation.

The Illusion of Openness

Open source is regarded as the great hope of AI democratisation. And it is true: the landscape has changed dramatically. In the spring of 2026, open models rival closed ones across numerous benchmarks.

[Source: Red Hat Developer, January 2026; Blue Headline, March 2026; Contabo, March 2026]

But 'open source' is not the same as open. The differences in practice are substantial.

Meta's Llama carries a licence that permits commercial use — provided you have fewer than 700 million monthly active users. This sounds generous, but it creates a ceiling: any competitor capable of threatening Meta automatically violates the licence.

DeepSeek uses an MIT licence — the most open available. But the training data, the fine-tuning methods, and parts of the architecture remain undocumented. And every output passes through censorship filters aligned with the CCP line.

[Source: Bruegel, 2025; Contabo, March 2026; Cultura Analysis, January 2026]

Mistral, the French company that represents Europe's strongest AI hope, publishes its models under Apache 2.0 — among the most open licences available. But its servers run in the clouds of Amazon and Microsoft.

[Source: Mistral AI; IntuitionLabs, November 2025; Red Hat Developer, January 2026]

And even when the models are open: the infrastructure on which they run is not. Cloud capacity is concentrated in three providers — AWS, Microsoft Azure, Google Cloud. Whoever controls the infrastructure controls the conditions of 'openness.'

What Public Infrastructure Could Mean

Imagine there were a public AI system — similar to the way public libraries, road networks, or the postal system exist. An infrastructure owned by no company, controlled by no state, governed democratically — transparently, accountably, for the benefit of all.

This sounds unrealistic. But approaches exist. France's President Macron announced in 2025 a 'historic partnership' with Nvidia: the Mistral Cloud project is to provide European researchers and companies with access to sovereign computing infrastructure.

[Source: IntuitionLabs/Mistral Large 3 Report, November 2025; VivaTech 2025]

Taken further, a 'public stack' could look like this: the base layer — computing capacity, foundation models, standardised datasets — is publicly provided. On top of it, private and state actors build. Access is open, conditions transparent, governance democratic.

The model is not without precedent. The internet itself began as public infrastructure — funded by the US Department of Defense, governed through open protocols, accessible to everyone. Only later came privatisation — and with it, concentration.

Elinor Ostrom showed in her research on commons that collectively managed resources — forests, fisheries, irrigation systems — need be neither privatised nor nationalised to function efficiently. They need rules, transparency, and the participation of those affected.

[Elinor Ostrom, Governing the Commons, 1990]

What Is at Stake — Three Scenarios

It helps to imagine the three architectures not as abstractions but as experiences. As the way a concrete person in a concrete city lives a concrete day.

In the first scenario — private infrastructure — a teacher in Dresden opens her tablet in the morning. Lesson preparation runs through an AI system for which the school district pays a licence. The system works well — but the topics are determined by the provider. When the licence becomes more expensive, the district must either pay or return to paper.

In the second scenario — state control — a programmer in Shenzhen uses a powerful, free AI model provided by the state. It works brilliantly — for everything the state permits. When she tries to research the history of protests, the model returns a blank screen.

In the third scenario — shared infrastructure — a social entrepreneur in Nairobi works with a publicly available AI system trained on African languages. He develops an application for smallholder farmers — without licence fees, without dependency, with full control over the data.

None of these scenarios is utopian or dystopian. Each has strengths and weaknesses. The first is efficient but creates dependency. The second is powerful but constrains freedom. The third is most equitable — but most complex to implement.

The question is not which model is perfect. The question is which we choose — and whether we choose consciously at all.

The Choice Being Made Right Now

This choice is not abstract. It is being made right now — in laws, in investments, in user decisions, in the way we talk about technology.

When a state rolls back AI regulation to accelerate innovation, it chooses the first model. When another state builds a sovereign computing platform, it chooses the second. And when European researchers build open models on public infrastructure, they are testing the third.

Karl Polanyi wrote that the great transformation of the 19th century consisted in re-embedding the market in society — through rules, institutions, social safety nets. The transformation of the 21st century consists in embedding AI — before it embeds us.

[Karl Polanyi, The Great Transformation, 1944]

Whether this succeeds depends not on the technology. It depends on us — on our capacity to think politically where technical thinking dominates. On our willingness to ask uncomfortable questions — and to accept that the answers will not be simple.

The final substantive chapter of this book asks the most personal of these questions: what does it mean to be human when machines can think? Not for states, not for economies — but for each of us.

Being Human in the Age of Machines

Identity, meaning, and the question that changes everything

The preceding chapters have described a world: power shifts, geopolitical rivalries, economic upheavals, synthetic publics, surveillance, and democracy under pressure.

But at the end, there is no number. At the end, there is a person. In the morning, before the mirror. With the question: who am I, if everything is changing?

This chapter is about that question. Not as psychological self-help, not as consolation — but as the question that gathers all the other questions of this book and poses it to each of us personally.

The Circle That Closes

The MindShift series began with a question about consciousness. 'You Are the Universe' asked: how does reality come into being? The answer was: not independently of us. We are not observers of the world — we are its participants.

The second book, 'Lost in Space,' deepened this insight in an existential direction: the search for belonging in a boundless world. The feeling of being able to go anywhere and yet belonging nowhere.

The third book, 'What Does AI Do to Us?,' opened the social dimension: how algorithmic systems shape not only our tools but our self-images, our relationships, our sense of normality.

The fourth book, 'Professional Identity in the AI Age,' made it concrete: what happens to self-worth when the work that defined us is done by machines?

The fifth book, 'The Quiet Shift,' examined the architecture of our perception: how feeds, algorithms, and synthetic realities determine what we see, what we think, and what we accept as the norm.

This sixth book has directed its gaze at the world as a whole: power, geopolitics, economy, war, democracy. But now, at the end, it returns — to the individual person. To the question that stood at the beginning.

What Machines Cannot Do

There is a widespread fear that can be summarised in a single sentence: if machines can think, what is left for us? If AI analyses, synthesises, optimises, creates — where is there still room for the human?

The fear is understandable. But it rests on a misunderstanding. It presupposes that being human consists in what we do. That our worth lies in our productivity.

The MindShift series has developed a different answer across five books: being human lies not in doing. It lies in being. In consciousness. In the capacity to feel, to doubt, to connect, to decide, to make meaning.

Hannah Arendt distinguished in *Vita Activa* three forms of human activity: labour (the necessary), work (the durable), and action (the free). Labour secures survival. Work builds the world. But only action — the beginning of something new and unpredictable — makes us human.

[Hannah Arendt, The Human Condition (Vita Activa), 1958]

AI can labour. It can work. But can it act — in the Arendtian sense? Can it begin something that is not derivable from its data? Can it engage in a deed that cannot be explained by optimisation?

The honest answer is: as of March 2026, it cannot. Whether it ever will is an open question. But even if it could one day — would our own action thereby lose its value? Does the beauty of a human deed diminish because a machine can imitate it?

Meaning Beyond Performance

Marie Jahoda studied the unemployed of Marienthal in the 1930s — an Austrian community whose factory had closed. What she found was not merely material deprivation. She found the dissolution of identity: without work, people lost structure, social bonds, self-respect, a sense of purpose.

[Marie Jahoda, The Unemployed of Marienthal, 1933]

Ninety years later, the pattern repeats — not as sudden factory closure but as creeping devaluation. Orders diminish. The period during which skills remain relevant shrinks. The confidence that one can hold on erodes.

Psychological research in recent years shows what happens. Studies by the University of Michigan, DIW Berlin, and the University of Konstanz from 2025 find: people who assess their work as threatened by AI display significantly lower life satisfaction, higher anxiety, and more frequent symptoms of depression.

[Source: University of Michigan PSID, 2025; DIW Berlin, 2025; University of Konstanz AI Study, 2025]

But the same research also shows something else: growth. The American Psychological Association reports that up to 70 percent of people who experience an identity crisis subsequently display 'post-traumatic growth' — a broadening of perspective, a deepening of relationships, a new understanding of their own values.

[Source: APA, Post-Traumatic Growth Studies, 2025]

The loss is real. And it is painful. But it is not the end. It is the space in which something new can emerge — if the conditions are right. If there is security, support, and time.

The Five Sources

Research in recent years — and the experience of many people already living through the transition — points to five sources of meaning when the familiar supports collapse.

Connection. The Harvard Grant Study — the longest longitudinal study in the world — has shown for decades: close relationships are the strongest predictor of long-term well-being. Not success, not wealth, not status — but the quality of human bonds.

Creativity. Not creativity as service but as expression. Drawing, writing, making music, building, gardening — not to sell, but because the process itself bestows meaning. Csikszentmihalyi called this 'flow' — the state of total immersion in which time vanishes and only the doing remains.

[Mihaly Csikszentmihalyi, Flow, 1990; updated studies 2025]

Contribution. Doing something for others — not from duty but from the need to belong to something larger. Studies of prosocial behaviour consistently show: people who regularly do something for others are more satisfied with their lives — regardless of income.

Presence. The capacity to be here — without a goal, without the pressure of productivity, without the constant question of whether one is doing enough. Mindfulness research shows: regular practice of presence reduces stress, improves concentration, and strengthens the sense of meaning.

Wonder. The capacity to be surprised by the world. By a sunset, a child's laughter, an insight one did not expect. Hartmut Rosa called this 'resonance' — the moment when the world touches us and we touch it back.

[Hartmut Rosa, Resonance, 2016]

Homo Ludens

Johan Huizinga coined the concept of Homo Ludens in 1938 — the playing human. He argued that play is older than culture, older than work, older than institutions. That people find their deepest essence not in labour but in free, self-purposeful creation — in play, ritual, art, exploration.

[Johan Huizinga, Homo Ludens, 1938]

If AI liberates us from labour — not completely, not immediately, but gradually — then the question arises: what takes its place? The answer of Homo Ludens is not in inactivity. It is in a different kind of activity.

What remains is something deeper: the possibility of being Homo Ludens. Not in the sense of entertainment or distraction. But in the sense of free creation, free exploration, free being — without the pressure that every action must be productive.

This is not a utopia. It is a description of what many people already experience — in the moments when they are most alive. In the moments when they are not working but playing, creating, wondering.

The Question That Remains

At the end of this chapter — and of this book — there is no answer. There is a question.

Not: what will AI do to us? But: what will we do with the freedom that arises when machines relieve us of necessity?

Will we fear it — and flee into new forms of busyness, just so as not to be alone with ourselves? Will we squander it — in distraction, consumption, boredom? Or will we use it — for what has always made us human: connection, creativity, contribution, presence, wonder?

The answer is not in this book. It is in every person who asks themselves this question — honestly, openly, without the convenient excuse that technology will sort everything out.

Technology sorts out nothing. It merely sharpens the question.

And the question is: who do we want to be?

12

Epilogue

Letter to a World in Transition

Dear reader,

I am writing these lines in March 2026, in my flat between Dresden and Izmir, between two worlds, both in motion. Through the window I see a city that does not know what it will look like in ten years. And neither do I.

But I have just written eleven chapters in which I tried to look closely. And what I have seen will not leave me.

I have seen that a handful of companies exerts more influence on our daily lives than most governments. That machines we do not understand make decisions about our work, our health, our security. That the line between reality and synthesis is blurring — and that most people do not notice.

I have seen that the economy produces more than ever and yet more and more people feel they cannot keep up. That the speed at which the world changes exceeds the speed at which we can adapt. That democracy — the form of coexistence we trusted — is under a pressure it has not experienced in decades.

And I have seen something that troubles me more than anything else: how quietly all of this is happening. No bang. No visible rupture. Only a slow shifting of the coordinates — and the feeling that the ground beneath our feet is no longer what it was.

Invisible Shift.

But I have also seen something else.

I have seen people who refuse to close their eyes. Researchers who publish uncomfortable truths. Politicians who push regulation forward even when the pressure is immense. Activists who create alternatives — small, imperfect, but real.

I have seen that pilot projects work — in Germany, in Catalonia, in Finland, in the Marshall Islands. That people who receive security do not become more passive — but freer.

And I have seen — in conversations, in messages, in the eyes of people who read these books — that the longing for depth has not vanished. That people do not want to be optimised. They want to be alive.

This book is the sixth in the MindShift series. It is the one with the greatest altitude — and perhaps the one that was hardest for me to write. Because it was necessary to call things by their names, without hiding behind either alarmism or optimism. Because the truth — in the middle — is uncomfortable and demands courage.

I do not know what the world will look like when you read these lines. Perhaps things have changed that I could not foresee. Perhaps the book is outdated in its details — the numbers, the laws, the constellations of power. The shelf life of facts is limited.

But the shelf life of the question it asks is not. For it will remain, whatever happens:

What kind of world do we want to build — and who do we want to be in it?

This question cannot be delegated. Not to machines, not to governments, not to markets, not to experts. It belongs to each individual. And it deserves to be asked — honestly, unhurriedly, and with the readiness to accept that the answer will not be simple.

Take the time. Look closely. Talk to the people who matter to you. Ask the questions that are uncomfortable. And trust that your capacity for judgement is your greatest strength. No algorithm can replace it.

The world is in transition. That is unsettling. But it also means: nothing is decided. Everything is still possible. And the power to shape it lies not in a silicon chip or in a data centre.

It lies in you.

Dresden / Izmir, March 2026

Hakan Özgür

www.mindshiftme.de