

#1 *New York Times* bestselling author of SAPIENS

# Yuval Noah Harari



# Nexus

A Brief History of Information Networks  
from the Stone Age to AI

BY YUVAL NOAH HARARI

*Sapiens*

*Homo Deus*

*21 Lessons for the 21st Century*

*Unstoppable Us*

*Nexus*

# Nexus

---

A BRIEF HISTORY OF INFORMATION NETWORKS  
FROM THE STONE AGE TO AI

Yuval Noah Harari



RANDOM HOUSE  
NEW YORK

Copyright © 2024 by Yuval Noah Harari

All rights reserved.

Published in the United States by Random House, an imprint and division of Penguin  
Random House LLC, New York.

RANDOM HOUSE and the HOUSE colophon are registered trademarks of Penguin Random  
House LLC.

Published in the United Kingdom by Fern Press, an imprint of Vintage, a division of  
Penguin Random House UK.

Library of Congress Cataloging-in-Publication Data

Names: Harari, Yuval N., author.

Title: Nexus: a brief history of information networks from the Stone Age to AI / Yuval  
Noah Harari.

Description: First edition. | New York: Random House, [2024] | Includes bibliographical  
references and index.

Identifiers: LCCN 2024011713 (print) | LCCN 2024011714 (ebook) | ISBN  
9780593734223 (hardcover) | ISBN 9780593736814 | ISBN 9780593734247 (ebook)

Subjects: LCSH: Information behavior—History. | Information networks—History. |  
Information technology—History.

Classification: LCC ZA3075 .H375 2024 (print) | LCC ZA3075 (ebook) | DDC 001.09—  
dc23/eng/20240614

LC record available at <https://lcn.loc.gov/2024011713>

LC ebook record available at <https://lcn.loc.gov/2024011714>

International edition ISBN 9780593736814

Ebook ISBN 9780593734247

[randomhousebooks.com](https://randomhousebooks.com)

*Book design by Caroline Cunningham, adapted for ebook*

Cover design: © Suzanne Dean

Cover illustration: © Hector the carrier pigeon of Emperor Napoleon III, reproduced by  
courtesy of Tallandier/Bridgeman Images

ep\_prh\_7.0\_148118692\_c0\_r0

# Contents

DEDICATION

PROLOGUE

## PART I: Human Networks

CHAPTER 1: What Is Information?

CHAPTER 2: Stories: Unlimited Connections

CHAPTER 3: Documents: The Bite of the Paper Tigers

CHAPTER 4: Errors: The Fantasy of Infallibility

CHAPTER 5: Decisions: A Brief History of Democracy and  
Totalitarianism

## PART II: The Inorganic Network

CHAPTER 6: The New Members: How Computers Are Different  
from Printing Presses

CHAPTER 7: Relentless: The Network Is Always On

CHAPTER 8: Fallible: The Network Is Often Wrong

## PART III: Computer Politics

CHAPTER 9: Democracies: Can We Still Hold a Conversation?

CHAPTER 10: Totalitarianism: All Power to the Algorithms?

CHAPTER 11: The Silicon Curtain: Global Empire or Global Split?

EPILOGUE

ACKNOWLEDGMENTS

NOTES

INDEX

ABOUT THE AUTHOR

*To Itzik with love, and to all who love wisdom.*

*On a path of a thousand dreams, we are looking for reality.*



# Prologue

---

We have named our species *Homo sapiens*—the wise human. But it is debatable how well we have lived up to the name.

Over the last 100,000 years, we Sapiens have certainly accumulated enormous power. Just listing all our discoveries, inventions, and conquests would fill volumes. But power isn't wisdom, and after 100,000 years of discoveries, inventions, and conquests humanity has pushed itself into an existential crisis. We are on the verge of ecological collapse, caused by the misuse of our own power. We are also busy creating new technologies like artificial intelligence (AI) that have the potential to escape our control and enslave or annihilate us. Yet instead of our species uniting to deal with these existential challenges, international tensions are rising, global cooperation is becoming more difficult, countries are stockpiling doomsday weapons, and a new world war does not seem impossible.

If we Sapiens are so wise, why are we so self-destructive?

At a deeper level, although we have accumulated so much information about everything from DNA molecules to distant galaxies, it doesn't seem that all this information has given us answers to the big questions of life: Who are we? What should we aspire to? What is a good life, and how should we live it? Despite the stupendous amounts of information at our disposal, we are as susceptible as our ancient ancestors to fantasy and delusion. Nazism and Stalinism are but two recent examples of the mass insanity that occasionally engulfs even modern societies. Nobody disputes

that humans today have a lot more information and power than in the Stone Age, but it is far from certain that we understand ourselves and our role in the universe much better.

Why are we so good at accumulating more information and power, but far less successful at acquiring wisdom? Throughout history many traditions have believed that some fatal flaw in our nature tempts us to pursue powers we don't know how to handle. The Greek myth of Phaethon told of a boy who discovers that he is the son of Helios, the sun god. Wishing to prove his divine origin, Phaethon demands the privilege of driving the chariot of the sun. Helios warns Phaethon that no human can control the celestial horses that pull the solar chariot. But Phaethon insists, until the sun god relents. After rising proudly in the sky, Phaethon indeed loses control of the chariot. The sun veers off course, scorching all vegetation, killing numerous beings, and threatening to burn the earth itself. Zeus intervenes and strikes Phaethon with a thunderbolt. The conceited human drops from the sky like a falling star, himself on fire. The gods reassert control of the sky and save the world.

Two thousand years later, when the Industrial Revolution was making its first steps and machines began replacing humans in numerous tasks, Johann Wolfgang von Goethe published a similar cautionary tale titled "The Sorcerer's Apprentice." Goethe's poem (later popularized as a Walt Disney animation starring Mickey Mouse) tells of an old sorcerer who leaves a young apprentice in charge of his workshop and gives him some chores to tend to while he is gone, like fetching water from the river. The apprentice decides to make things easier for himself and, using one of the sorcerer's spells, enchants a broom to fetch the water for him. But the apprentice doesn't know how to stop the broom, which relentlessly fetches more and more water, threatening to flood the workshop. In panic, the apprentice cuts the enchanted broom in two with an ax, only to see each half become another broom. Now *two* enchanted brooms are inundating the workshop with water. When the old sorcerer returns, the apprentice pleads for help: "The spirits that I summoned, I now cannot rid myself of again." The sorcerer immediately breaks the spell and stops the flood. The lesson to the

apprentice—and to humanity—is clear: never summon powers you cannot control.

What do the cautionary fables of the apprentice and of Phaethon tell us in the twenty-first century? We humans have obviously refused to heed their warnings. We have already driven the earth's climate out of balance and have summoned billions of enchanted brooms, drones, chatbots, and other algorithmic spirits that may escape our control and unleash a flood of unintended consequences.

What should we do, then? The fables offer no answers, other than to wait for some god or sorcerer to save us. This, of course, is an extremely dangerous message. It encourages people to abdicate responsibility and put their faith in gods and sorcerers instead. Even worse, it fails to appreciate that gods and sorcerers are themselves a human invention—just like chariots, brooms, and algorithms. The tendency to create powerful things with unintended consequences started not with the invention of the steam engine or AI but with the invention of religion. Prophets and theologians have summoned powerful spirits that were supposed to bring love and joy but occasionally ended up flooding the world with blood.

The Phaethon myth and Goethe's poem fail to provide useful advice because they misconstrue the way humans gain power. In both fables, a single human acquires enormous power, but is then corrupted by hubris and greed. The conclusion is that our flawed individual psychology makes us abuse power. What this crude analysis misses is that human power is never the outcome of individual initiative. Power always stems from cooperation between large numbers of humans.

Accordingly, it isn't our individual psychology that causes us to abuse power. After all, alongside greed, hubris, and cruelty, humans are also capable of love, compassion, humility, and joy. True, among the worst members of our species, greed and cruelty reign supreme and lead bad actors to abuse power. But why would human societies choose to entrust power to their worst members? Most Germans in 1933, for example, were not psychopaths. So why did they vote for Hitler?

Our tendency to summon powers we cannot control stems not from individual psychology but from the unique way our species cooperates in large numbers. The main argument of this book is that humankind gains enormous power by building large networks of cooperation, but the way these networks are built predisposes us to use that power unwisely. Our problem, then, is a network problem.


Even more specifically, it is an information problem. Information is the glue that holds networks together. But for tens of thousands of years, Sapiens built and maintained large networks by inventing and spreading fictions, fantasies, and mass delusions—about gods, about enchanted broomsticks, about AI, and about a great many other things. While each individual human is typically interested in knowing the truth about themselves and the world, large networks bind members and create order by relying on fictions and fantasies. That's how we got, for example, to Nazism and Stalinism. These were exceptionally powerful networks, held together by exceptionally deluded ideas. As George Orwell famously put it, ignorance is strength.

The fact that the Nazi and Stalinist regimes were founded on cruel fantasies and shameless lies did *not* make them historically exceptional, nor did it preordain them to collapse. Nazism and Stalinism were two of the strongest networks humans ever created. In late 1941 and early 1942, the Axis powers came within reach of winning World War II. Stalin eventually emerged as the victor of that war,<sup>[1]</sup> and in the 1950s and 1960s he and his heirs also had a reasonable chance of winning the Cold War. By the 1990s liberal democracies had gained the upper hand, but this now seems like a temporary victory. In the twenty-first century, some new totalitarian regime may well succeed where Hitler and Stalin failed, creating an all-powerful network that could prevent future generations from even attempting to expose its lies and fictions. We should not assume that delusional networks are doomed to failure. If we want to prevent their triumph, we will have to do the hard work ourselves.

## THE NAIVE VIEW OF INFORMATION

It is difficult to appreciate the strength of delusional networks because of a broader misunderstanding about how big information networks—whether delusional or not—operate. This misunderstanding is encapsulated in something I call “the naive view of information.” While fables like the myth of Phaethon and “The Sorcerer’s Apprentice” present an overly pessimistic view of individual human psychology, the naive view of information disseminates an overly optimistic view of large-scale human networks.

The naive view argues that by gathering and processing much more information than individuals can, big networks achieve a better understanding of medicine, physics, economics, and numerous other fields, which makes the network not only powerful but also wise. For example, by gathering information on pathogens, pharmaceutical companies and health-care services can determine the true causes of many diseases, which enables them to develop more effective medicines and to make wiser decisions about their usage. This view posits that in sufficient quantities information leads to truth, and truth in turn leads to both power and wisdom. Ignorance, in contrast, seems to lead nowhere. While delusional or deceitful networks might occasionally arise in moments of historical crisis, in the long term they are bound to lose to more clear-sighted and honest rivals. A health-care service that ignores information about pathogens, or a pharmaceutical giant that deliberately spreads disinformation, will ultimately lose out to competitors that make wiser use of information. The naive view thus implies that delusional networks must be aberrations and that big networks can usually be trusted to handle power wisely.



### The naive view of information

Of course, the naive view acknowledges that many things can go wrong on the path from information to truth. We might make honest mistakes in gathering and processing the information. Malicious actors motivated by greed or hate might hide important facts or try to deceive us. As a result, information sometimes leads to error rather than truth. For example, partial information, faulty analysis, or a disinformation campaign might lead even experts to misidentify the true cause of a particular disease.

However, the naive view assumes that the antidote to most problems we encounter in gathering and processing information is gathering and processing even more information. While we are never completely safe from error, in most cases more information means greater accuracy. A single doctor wishing to identify the cause of an epidemic by examining a single patient is less likely to succeed than thousands of doctors gathering data on millions of patients. And if the doctors themselves conspire to hide the truth, making medical information more freely available to the public and to investigative journalists will eventually reveal the scam. According to this view, the bigger the information network, the closer it must be to the truth.

Naturally, even if we analyze information accurately and discover important truths, this does not guarantee we will use the resulting capabilities wisely. Wisdom is commonly understood to mean “making right decisions,” but what “right” means depends on value judgments that differ among diverse people, cultures, and ideologies. Scientists who discover a new pathogen may develop a vaccine to protect people. But if the scientists—or their political overlords—believe in a racist ideology that advocates that some races are inferior and should be exterminated, the new

medical knowledge might be used to develop a biological weapon that kills millions.

In this case, too, the naive view of information holds that additional information offers at least a partial remedy. The naive view thinks that disagreements about values turn out on closer inspection to be the fault of either the lack of information or deliberate disinformation. According to this view, racists are ill-informed people who just don't know the facts of biology and history. They think that "race" is a valid biological category, and they have been brainwashed by bogus conspiracy theories. The remedy to racism is therefore to provide people with more biological and historical facts. It may take time, but in a free market of information sooner or later truth will prevail.

The naive view is of course more nuanced and thoughtful than can be explained in a few paragraphs, but its core tenet is that information is an essentially good thing, and the more we have of it, the better. Given enough information and enough time, we are bound to discover the truth about things ranging from viral infections to racist biases, thereby developing not only our power but also the wisdom necessary to use that power well.

This naive view justifies the pursuit of ever more powerful information technologies and has been the semiofficial ideology of the computer age and the internet. In June 1989, a few months before the fall of the Berlin Wall and of the Iron Curtain, Ronald Reagan declared that "the Goliath of totalitarian control will rapidly be brought down by the David of the microchip" and that "the biggest of Big Brothers is increasingly helpless against communications technology.... Information is the oxygen of the modern age.... It seeps through the walls topped with barbed wire. It wafts across the electrified, booby-trapped borders. Breezes of electronic beams blow through the Iron Curtain as if it was lace."<sup>[2]</sup> In November 2009, Barack Obama spoke in the same spirit on a visit to Shanghai, telling his Chinese hosts, "I am a big believer in technology and I'm a big believer in openness when it comes to the flow of information. I think that the more freely information flows, the stronger the society becomes."<sup>[3]</sup>

Entrepreneurs and corporations have often expressed similarly rosy views of information technology. Already in 1858 an editorial in *The New Englander* about the invention of the telegraph stated, “It is impossible that old prejudices and hostilities should longer exist, while such an instrument has been created for an exchange of thought between all the nations of the earth.”<sup>[4]</sup> Nearly two centuries and two world wars later, Mark Zuckerberg said that Facebook’s goal “is to help people to share more in order to make the world more open and to help promote understanding between people.”<sup>[5]</sup>

In his 2024 book, *The Singularity Is Nearer*, the eminent futurologist and entrepreneur Ray Kurzweil surveys the history of information technology and concludes that “the reality is that nearly every aspect of life is getting progressively better as a result of exponentially improving technology.” Looking back at the grand sweep of human history, he cites examples like the invention of the printing press to argue that by its very nature information technology tends to spawn “a virtuous circle advancing nearly every aspect of human well-being, including literacy, education, wealth, sanitation, health, democratization and reduction in violence.”<sup>[6]</sup>

The naive view of information is perhaps most succinctly captured in Google’s mission statement “to organize the world’s information and make it universally accessible and useful.” Google’s answer to Goethe’s warnings is that while a single apprentice pilfering his master’s secret spell book is likely to cause disaster, when a lot of apprentices are given free access to all the world’s information, they will not only create useful enchanted brooms but also learn to handle them wisely.

## GOOGLE VERSUS GOETHE

It must be stressed that there are numerous cases in which having more information has indeed enabled humans to understand the world better and to make wiser use of their power. Consider, for example, the dramatic reduction in child mortality. Johann Wolfgang von Goethe was the eldest of



seven siblings, but only he and his sister Cornelia got to celebrate their seventh birthday. Disease carried off their brother Hermann Jacob at age six, their sister Catharina Elisabeth at age four, their sister Johanna Maria at age two, their brother Georg Adolf at age eight months, and a fifth, unnamed brother was stillborn. Cornelia then died from disease at twenty-six, leaving Johann Wolfgang as the sole survivor from their family.<sup>[7]</sup>

Johann Wolfgang von Goethe went on to have five children of his own, of whom all but the eldest son—August—died within two weeks of their birth. In all probability the cause was incompatibility between the blood groups of Goethe and his wife, Christiane, which after the first successful pregnancy led the mother to develop antibodies to the fetal blood. This condition, known as rhesus disease, is nowadays treated so effectively that the mortality rate is less than 2 percent, but in the 1790s it had an average mortality rate of 50 percent, and for Goethe's four younger children it was a death sentence.<sup>[8]</sup>

Altogether in Goethe's family—a well-to-do German family in the late eighteenth century—the child survival rate was an abysmal 25 percent. Only three out of twelve children reached adulthood. This horrendous statistic was not exceptional. Around the time Goethe wrote “The Sorcerer's Apprentice” in 1797, it is estimated that only about 50 percent of German children reached age fifteen,<sup>[9]</sup> and the same was probably true in most other parts of the world.<sup>[10]</sup> By 2020, 95.6 percent of children worldwide lived beyond their fifteenth birthday,<sup>[11]</sup> and in Germany that figure was 99.5 percent.<sup>[12]</sup> This momentous achievement would not have been possible without collecting, analyzing, and sharing massive amounts of medical data about things like blood groups. In this case, then, the naive view of information proved to be correct.

However, the naive view of information sees only part of the picture, and the history of the modern age was not just about reducing child mortality. In recent generations humanity has experienced the greatest increase ever in both the amount and the speed of our information production. Every smartphone contains more information than the ancient Library of Alexandria<sup>[13]</sup> and enables its owner to instantaneously connect to billions

of other people throughout the world. Yet with all this information circulating at breathtaking speeds, humanity is closer than ever to annihilating itself.

Despite—or perhaps because of—our hoard of data, we are continuing to spew greenhouse gases into the atmosphere, pollute rivers and oceans, cut down forests, destroy entire habitats, drive countless species to extinction, and jeopardize the ecological foundations of our own species. We are also producing ever more powerful weapons of mass destruction, from thermonuclear bombs to doomsday viruses. Our leaders don't lack information about these dangers, yet instead of collaborating to find solutions, they are edging closer to a global war.

Would having even more information make things better—or worse? We will soon find out. Numerous corporations and governments are in a race to develop the most powerful information technology in history—AI. Some leading entrepreneurs, like the American investor Marc Andreessen, believe that AI will finally solve all of humanity's problems. On June 6, 2023, Andreessen published an essay titled “Why AI Will Save the World,” peppered with bold statements like “I am here to bring the good news: AI will not destroy the world, and in fact may save it” and “AI can make everything we care about better.” He concluded, “The development and proliferation of AI—far from a risk that we should fear—is a moral obligation that we have to ourselves, to our children, and to our future.”<sup>[14]</sup>

Ray Kurzweil concurs, arguing in *The Singularity Is Nearer* that “AI is the pivotal technology that will allow us to meet the pressing challenges that confront us, including overcoming disease, poverty, environmental degradation, and all of our human frailties. We have a moral imperative to realize this promise of new technologies.” Kurzweil is keenly aware of the technology's potential perils, and analyzes them at length, but believes they could be mitigated successfully.<sup>[15]</sup>

Others are more skeptical. Not only philosophers and social scientists but also many leading AI experts and entrepreneurs like Yoshua Bengio, Geoffrey Hinton, Sam Altman, Elon Musk, and Mustafa Suleyman have warned the public that AI could destroy our civilization.<sup>[16]</sup> A 2024 article

co-authored by Bengio, Hinton, and numerous other experts noted that “unchecked AI advancement could culminate in a large-scale loss of life and the biosphere, and the marginalization or even extinction of humanity.”<sup>[17]</sup> In a 2023 survey of 2,778 AI researchers, more than a third gave at least a 10 percent chance to advanced AI leading to outcomes as bad as human extinction.<sup>[18]</sup> In 2023 close to thirty governments—including those of China, the United States, and the U.K.—signed the Bletchley Declaration on AI, which acknowledged that “there is potential for serious, even catastrophic, harm, either deliberate or unintentional, stemming from the most significant capabilities of these AI models.”<sup>[19]</sup> By using such apocalyptic terms, experts and governments have no wish to conjure a Hollywood image of rebellious robots running in the streets and shooting people. Such a scenario is unlikely, and it merely distracts people from the real dangers. Rather, experts warn about two other scenarios.

First, the power of AI could supercharge existing human conflicts, dividing humanity against itself. Just as in the twentieth century the Iron Curtain divided the rival powers in the Cold War, so in the twenty-first century the Silicon Curtain—made of silicon chips and computer codes rather than barbed wire—might come to divide rival powers in a new global conflict. Because the AI arms race will produce ever more destructive weapons, even a small spark might ignite a cataclysmic conflagration.

Second, the Silicon Curtain might come to divide not one group of humans from another but rather all humans from our new AI overlords. No matter where we live, we might find ourselves cocooned by a web of unfathomable algorithms that manage our lives, reshape our politics and culture, and even reengineer our bodies and minds—while we can no longer comprehend the forces that control us, let alone stop them. If a twenty-first-century totalitarian network succeeds in conquering the world, it may be run by nonhuman intelligence, rather than by a human dictator. People who single out China, Russia, or a post-democratic United States as their main source for totalitarian nightmares misunderstand the danger. In fact, Chinese, Russians, Americans, and all other humans are together threatened by the totalitarian potential of nonhuman intelligence.

Given the magnitude of the danger, AI should be of interest to all human beings. While not everyone can become an AI expert, we should all keep in mind that AI is the first technology in history that can make decisions and create new ideas by itself. All previous human inventions have empowered humans, because no matter how powerful the new tool was, the decisions about its usage remained in our hands. Knives and bombs do not themselves decide whom to kill. They are dumb tools, lacking the intelligence necessary to process information and make independent decisions. In contrast, AI can process information by itself, and thereby replace humans in decision making. AI isn't a tool—it's an agent.

Its mastery of information also enables AI to independently generate new ideas, in fields ranging from music to medicine. Gramophones played our music, and microscopes revealed the secrets of our cells, but gramophones couldn't compose new symphonies, and microscopes couldn't synthesize new drugs. AI is already capable of producing art and making scientific discoveries by itself. In the next few decades, it will likely gain the ability even to create new life-forms, either by writing genetic code or by inventing an inorganic code animating inorganic entities.

Even at the present moment, in the embryonic stage of the AI revolution, computers already make decisions about us—whether to give us a mortgage, to hire us for a job, to send us to prison. This trend will only increase and accelerate, making it more difficult to understand our own lives. Can we trust computer algorithms to make wise decisions and create a better world? That's a much bigger gamble than trusting an enchanted broom to fetch water. And it is more than just human lives we are gambling on. AI could alter the course not just of our species' history but of the evolution of all life-forms.

## WEAPONIZING INFORMATION

In 2016, I published *Homo Deus*, a book that highlighted some of the dangers posed to humanity by the new information technologies. That book

argued that the real hero of history has always been information, rather than *Homo sapiens*, and that scientists increasingly understand not just history but also biology, politics, and economics in terms of information flows. Animals, states, and markets are all information networks, absorbing data from the environment, making decisions, and releasing data back. The book warned that while we hope better information technology will give us health, happiness, and power, it may actually take power away from us and destroy both our physical and our mental health. *Homo Deus* hypothesized that if humans aren't careful, we might dissolve within the torrent of information like a clump of earth within a gushing river, and that in the grand scheme of things humanity will turn out to have been just a ripple within the cosmic dataflow.

In the years since *Homo Deus* was published, the pace of change has only accelerated, and power has indeed been shifting from humans to algorithms. Many of the scenarios that sounded like science fiction in 2016—such as algorithms that can create art, masquerade as human beings, make crucial life decisions about us, and know more about us than we know about ourselves—are everyday realities in 2024.

Many other things have changed since 2016. The ecological crisis has intensified, international tensions have escalated, and a populist wave has undermined the cohesion of even the most robust democracies. Populism has also mounted a radical challenge to the naive view of information. Populist leaders such as Donald Trump and Jair Bolsonaro, and populist movements and conspiracy theories such as QAnon and the anti-vaxxers, have argued that all traditional institutions that gain authority by claiming to gather information and discover truth are simply lying. Bureaucrats, judges, doctors, mainstream journalists, and academic experts are elite cabals that have no interest in the truth and are deliberately spreading disinformation to gain power and privileges for themselves at the expense of “the people.” The rise of politicians like Trump and movements like QAnon has a specific political context, unique to the conditions of the United States in the late 2010s. But populism as an antiestablishment worldview long

predated Trump and is relevant to numerous other historical contexts now and in the future. In a nutshell, populism views information as a weapon.<sup>[20]</sup>

Information → Power

### The populist view of information

In its more extreme versions, populism posits that there is no objective truth at all and that everyone has “their own truth,” which they wield to vanquish rivals. According to this worldview, power is the only reality. All social interactions are power struggles, because humans are interested only in power. The claim to be interested in something else—like truth or justice—is nothing more than a ploy to gain power. Whenever and wherever populism succeeds in disseminating the view of information as a weapon, language itself is undermined. Nouns like “facts” and adjectives like “accurate” and “truthful” become elusive. Such words are not taken as pointing to a common objective reality. Rather, any talk of “facts” or “truth” is bound to prompt at least some people to ask, “Whose facts and whose truth are you referring to?”

It should be stressed that this power-focused and deeply skeptical view of information isn’t a new phenomenon and it wasn’t invented by anti-vaxxers, flat-earthers, Bolsonaristas, or Trump supporters. Similar views were propagated long before 2016, including by some of humanity’s brightest minds.<sup>[21]</sup> In the late twentieth century, for example, intellectuals from the radical left like Michel Foucault and Edward Said claimed that scientific institutions like clinics and universities are not pursuing timeless and objective truths but are instead using power to determine what counts as truth, in the service of capitalist and colonialist elites. These radical critiques occasionally went as far as arguing that “scientific facts” are nothing more than a capitalist or colonialist “discourse” and that people in power can never be really interested in truth and can never be trusted to recognize and correct their own mistakes.<sup>[22]</sup>

This particular line of radical leftist thinking goes back to Karl Marx, who argued in the mid-nineteenth century that power is the only reality, that information is a weapon, and that elites who claim to be serving truth and justice are in fact pursuing narrow class privileges. In the words of the 1848 *Communist Manifesto*, “The history of all hitherto existing societies is the history of class struggles. Freeman and slave, patrician and plebeian, lord and serf, guildmaster and journeyman, in a word, oppressor and oppressed stood in constant opposition to one another, carried on an uninterrupted, now hidden, now open, fight.” This binary interpretation of history implies that every human interaction is a power struggle between oppressors and oppressed. Accordingly, whenever anyone says anything, the question to ask isn’t, “What is being said? Is it true?” but rather, “Who is saying this? Whose privileges does it serve?”

Of course, right-wing populists such as Trump and Bolsonaro are unlikely to have read Foucault or Marx, and indeed present themselves as fiercely anti-Marxist. They also greatly differ from Marxists in their suggested policies in fields like taxation and welfare. But their basic view of society and of information is surprisingly Marxist, seeing all human interactions as a power struggle between oppressors and oppressed. For example, in his inaugural address in 2017 Trump announced that “a small group in our nation’s capital has reaped the rewards of government while the people have borne the cost.”<sup>[23]</sup> Such rhetoric is a staple of populism, which the political scientist Cas Mudde has described as an “ideology that considers society to be ultimately separated into two homogeneous and antagonistic groups, ‘the pure people’ versus ‘the corrupt elite.’”<sup>[24]</sup> Just as Marxists claimed that the media functions as a mouthpiece for the capitalist class, and that scientific institutions like universities spread disinformation in order to perpetuate capitalist control, populists accuse these same institutions of working to advance the interests of the “corrupt elites” at the expense of “the people.”

Present-day populists also suffer from the same incoherence that plagued radical antiestablishment movements in previous generations. If power is the only reality, and if information is just a weapon, what does it imply

about the populists themselves? Are they too interested only in power, and are they too lying to us to gain power?

Populists have sought to extricate themselves from this conundrum in two different ways. Some populist movements claim adherence to the ideals of modern science and to the traditions of skeptical empiricism. They tell people that indeed you should never trust any institutions or figures of authority—including self-proclaimed populist parties and politicians. Instead, you should “do your own research” and trust only what you can directly observe by yourself.<sup>[25]</sup> This radical empiricist position implies that while large-scale institutions like political parties, courts, newspapers, and universities can never be trusted, individuals who make the effort can still find the truth by themselves.

This approach may sound scientific and may appeal to free-spirited individuals, but it leaves open the question of how human communities can cooperate to build health-care systems or pass environmental regulations, which demand large-scale institutional organization. Is a single individual capable of doing all the necessary research to decide whether the earth’s climate is heating up and what should be done about it? How would a single person go about collecting climate data from throughout the world, not to mention obtaining reliable records from past centuries? Trusting only “my own research” may sound scientific, but in practice it amounts to believing that there is no objective truth. As we shall see in chapter 4, science is a collaborative institutional effort rather than a personal quest.

An alternative populist solution is to abandon the modern scientific ideal of finding the truth via “research” and instead go back to relying on divine revelation or mysticism. Traditional religions like Christianity, Islam, and Hinduism have typically characterized humans as untrustworthy power-hungry creatures who can access the truth only thanks to the intervention of a divine intelligence. In the 2010s and early 2020s populist parties from Brazil to Turkey and from the United States to India have aligned themselves with such traditional religions. They have expressed radical doubt about modern institutions while declaring complete faith in ancient scriptures. The populists claim that the articles you read in *The New York*



*Times* or in *Science* are just an elitist ploy to gain power, but what you read in the Bible, the Quran, or the Vedas is absolute truth.<sup>[26]</sup>

A variation on this theme calls on people to put their trust in charismatic leaders like Trump and Bolsonaro, who are depicted by their supporters either as the messengers of God<sup>[27]</sup> or as possessing a mystical bond with “the people.” While ordinary politicians lie to the people in order to gain power for themselves, the charismatic leader is the infallible mouthpiece of the people who exposes all the lies.<sup>[28]</sup> One of the recurrent paradoxes of populism is that it starts by warning us that all human elites are driven by a dangerous hunger for power, but often ends by entrusting all power to a single ambitious human.

We will explore populism at greater depth in chapter 5, but at this point it is important to note that populists are eroding trust in large-scale institutions and international cooperation just when humanity confronts the existential challenges of ecological collapse, global war, and out-of-control technology. Instead of trusting complex human institutions, populists give us the same advice as the Phaethon myth and “The Sorcerer’s Apprentice”: “Trust God or the great sorcerer to intervene and make everything right again.” If we take this advice, we’ll likely find ourselves in the short term under the thumb of the worst kind of power-hungry humans, and in the long term under the thumb of new AI overlords. Or we might find ourselves nowhere at all, as Earth becomes inhospitable for human life.

If we wish to avoid relinquishing power to a charismatic leader or an inscrutable AI, we must first gain a better understanding of what information is, how it helps to build human networks, and how it relates to truth and power. Populists are right to be suspicious of the naive view of information, but they are wrong to think that power is the only reality and that information is always a weapon. Information isn’t the raw material of truth, but it isn’t a mere weapon, either. There is enough space between these extremes for a more nuanced and hopeful view of human information networks and of our ability to handle power wisely. This book is dedicated to exploring that middle ground.

## THE ROAD AHEAD

The first part of this book surveys the historical development of human information networks. It doesn't attempt to present a comprehensive century-by-century account of information technologies like script, printing presses, and radio. Instead, by studying a few examples, it explores key dilemmas that people in all eras faced when trying to construct information networks, and it examines how different answers to these dilemmas shaped contrasting human societies. What we usually think of as ideological and political conflicts often turn out to be clashes between opposing types of information networks.

Part 1 begins by examining two principles that have been essential for large-scale human information networks: mythology and bureaucracy. Chapters 2 and 3 describe how large-scale information networks—from ancient kingdoms to present-day states—have relied on both mythmakers and bureaucrats. The stories of the Bible, for example, were essential for the Christian Church, but there would have been no Bible if church bureaucrats hadn't curated, edited, and disseminated these stories. A difficult dilemma for every human network is that mythmakers and bureaucrats tend to pull in different directions. Institutions and societies are often defined by the balance they manage to find between the conflicting needs of their mythmakers and their bureaucrats. The Christian Church itself split into rival churches, like the Catholic and Protestant churches, which struck different balances between mythology and bureaucracy.

Chapter 4 then focuses on the problem of erroneous information and on the benefits and drawbacks of maintaining self-correcting mechanisms, such as independent courts or peer-reviewed journals. The chapter contrasts institutions that relied on weak self-correcting mechanisms, like the Catholic Church, with institutions that developed strong self-correcting mechanisms, like scientific disciplines. Weak self-correcting mechanisms sometimes result in historical calamities like the early modern European witch hunts, while strong self-correcting mechanisms sometimes destabilize the network from within. Judged in terms of longevity, spread, and power,

the Catholic Church has been perhaps the most successful institution in human history, despite—or perhaps because of—the relative weakness of its self-correcting mechanisms.

After part 1 surveys the roles of mythology and bureaucracy, and the contrast between strong and weak self-correcting mechanisms, chapter 5 concludes the historical discussion by focusing on another contrast—between distributed and centralized information networks. Democratic systems allow information to flow freely along many independent channels, whereas totalitarian systems strive to concentrate information in one hub. Each choice has both advantages and shortcomings. Understanding political systems like the United States and the U.S.S.R. in terms of information flows can explain much about their differing trajectories.

This historical part of the book is crucial for understanding present-day developments and future scenarios. The rise of AI is arguably the biggest information revolution in history. But we cannot understand it unless we compare it with its predecessors. History isn't the study of the past; it is the study of change. History teaches us what remains the same, what changes, and how things change. This is as relevant to information revolutions as to every other kind of historical transformation. Thus, understanding the process through which the allegedly infallible Bible was canonized provides valuable insight about present-day claims for AI infallibility. Similarly, studying the early modern witch hunts and Stalin's collectivization offers stark warnings about what might go wrong as we give AIs greater control over twenty-first-century societies. A deep knowledge of history is also vital to understand what *is* new about AI, how it is fundamentally different from printing presses and radio sets, and in what specific ways an AI dictatorship could be very *unlike* anything we have seen before.

The book doesn't argue that studying the past enables us to predict the future. As emphasized repeatedly in the following pages, history is not deterministic, and the future will be shaped by the choices we all make in coming years. The whole point of writing this book is that by making informed choices, we can prevent the worst outcomes. If we cannot change the future, why waste time discussing it?

Building upon the historical survey in part 1, the book's second part—"The Inorganic Network"—examines the new information network we are creating today, focusing on the political implications of the rise of AI. Chapters 6–8 discuss recent examples from throughout the world—such as the role of social media algorithms in instigating ethnic violence in Myanmar in 2016–17—to explain in what ways AI is different from all previous information technologies. Examples are taken mostly from the 2010s rather than the 2020s, because we have gained a modicum of historical perspective on events of the 2010s.

Part 2 argues that we are creating an entirely new kind of information network, without pausing to reckon with its implications. It emphasizes the shift from organic to inorganic information networks. The Roman Empire, the Catholic Church, and the U.S.S.R. all relied on carbon-based brains to process information and make decisions. The silicon-based computers that dominate the new information network function in radically different ways. For better or worse, silicon chips are free from many of the limitations that organic biochemistry imposes on carbon neurons. Silicon chips can create spies that never sleep, financiers that never forget, and despots that never die. How will this change society, economics, and politics?

The third and final part of the book—"Computer Politics"—examines how different kinds of societies might deal with the threats and promises of the inorganic information network. Will carbon-based life-forms like us have a chance of understanding and controlling the new information network? As noted above, history isn't deterministic, and for at least a few more years we Sapiens still have the power to shape our future.

Accordingly, chapter 9 explores how democracies might deal with the inorganic network. How, for example, can flesh-and-blood politicians make financial decisions if the financial system is increasingly controlled by AI and the very meaning of money comes to depend on inscrutable algorithms? How can democracies maintain a public conversation about anything—be it finance or gender—if we can no longer know whether we are talking with another human or with a chatbot masquerading as a human?

Chapter 10 explores the potential impact of the inorganic network on totalitarianism. While dictators would be happy to get rid of all public conversations, they have their own fears of AI. Autocracies are based on terrorizing and censoring their own agents. But how can a human dictator terrorize an AI, censor its unfathomable processes, or prevent it from seizing power for itself?

Finally, chapter 11 explores how the new information network could influence the balance of power between democratic and totalitarian societies on the global level. Will AI tilt the balance decisively in favor of one camp? Will the world split into hostile blocs whose rivalry makes all of us easy prey for an out-of-control AI? Or can we unite in defense of our common interests?

But before we explore the past, present, and possible futures of information networks, we need to start with a deceptively simple question. What exactly is information?