

WITH NEW AND UPDATED MATERIAL

Yuval Noah Harari

'A deeply
important book'

MUSTAFA
SULEYMAN


'Tremendous...
Thought-provoking'

STEPHEN
FRY



'Provocative
and entrancing'

RORY
STEWART



Nexus

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

‘A historian whose arguments operate on the scale of millennia has managed to capture the zeitgeist perfectly . . . Harari’s narrative is engaging and his framing is strikingly original’

The Economist

‘Harari sticks the world together in a gleaming shape that inspires and excites’

Telegraph

‘A superb narrative writer’

Guardian

‘Engrossing . . . A diagnosis and a call to action’

Observer

‘[Harari] makes fun of people like myself who saw more information as always a good thing . . . I would basically say he’s right and I was wrong’

Bill Gates

‘A useful, well-informed primer . . . Wise and bold’

New York Times

‘One of our foremost big-picture thinkers about the grand sweep of humanity, history and the future’

TIME

‘*Nexus* is an important book and comes at an important time in human history’

Science

‘*Nexus* is ambitious, bold and, at times, unsettling . . . For anyone interested in the intersection of history, technology and power, Harari once again provokes deep thought’

The Conversation

‘Yuval Noah Harari is one of the most remarkable intellects of our generation – bold, original, erudite, provocative and entrancing. His latest book reimagines everything from literacy to AI and . . . all his books fundamentally shifts one’s view of the world’

Rory Stewart

‘Tremendous, thought-provoking and so very well reasoned. Harari gives us a vision of a rapidly approaching future that is at one and the same time thrilling and chilling. If there is one book that I would urge everyone to read – our political, corporate and cultural leaders most especially – it is *Nexus*’

Stephen Fry

‘Yuval Noah Harari has a unique ability to unite history’s finest details and its grandest megatrends in a single view. In this masterful and provocative new book, he makes a compelling case that information networks are – and always have been – the primary driving force shaping human societies. This deeply important book comes at a critical time as we all think through the implications of AI and automated content production’

Mustafa Suleyman

‘*Nexus* will challenge your core beliefs about technology and information while leaving you grateful for the experience’

Dr Joy Buolamwini

‘In *Nexus*, Harari uses information networks as the backbone of human evolution, offering profound insights into how technological progress reshapes our social structures and individual lives. His unique perspectives and deep understanding of the age of artificial intelligence is especially pertinent in the current wave of transformations driven by generative AI. The book urges us to reassess the relationship between humans and AI and how we can ensure that technological progress serves the greater good for humanity’

Kai-Fu Lee

‘[Harari] connects past and future in a way that’s both enlightening and a little eerie. *Nexus* isn’t just a critique of technology; it’s a reflection on how we’ve always been vulnerable to those who control information. It’s a wake-up call in the gentlest, most urgent way’

Yotam Ottolenghi

YUVAL NOAH HARARI

Yuval Noah Harari is a historian, philosopher and the global bestselling author of *Sapiens*; *Homo Deus*; *21 Lessons for the 21st Century*; the graphic adaptation series *Sapiens: A Graphic History*; and *Unstoppable Us*, his first series for children. His books have sold more than 50 million copies in 65 languages. He co-founded Sapienship – an international social-impact company focused on education and storytelling – with his husband, Itzik Yahav. Harari has a PhD from the University of Oxford and has been a lecturer at the Hebrew University of Jerusalem. He is a distinguished research fellow at the University of Cambridge’s Centre for the Study of Existential Risk.

Copyrighted Material

ALSO BY YUVAL NOAH HARARI

Sapiens

Homo Deus

21 Lessons for the 21st Century

Sapiens: A Graphic History

Unstoppable Us

Copyrighted Material

YUVAL NOAH HARARI

Nexus

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

Copyrighted Material
VINTAGE

1 3 5 7 9 10 8 6 4 2

Vintage is part of the Penguin Random House group of companies

Vintage, Penguin Random House UK, One Embassy Gardens,
8 Viaduct Gardens, London SW11 7BW

penguin.co.uk/vintage
global.penguinrandomhouse.com



Penguin
Random House
UK

First published in Vintage in 2025

First published in hardback by Fern Press in 2024

First published in the United States of America by Random House,
a division of Penguin Random House LLC, New York, in 2024

Copyright © Yuval Noah Harari 2024, 2025

The moral right of the author has been asserted

Book design by Caroline Cunningham

Penguin Random House values and supports copyright. Copyright fuels creativity, encourages diverse voices, promotes freedom of expression and supports a vibrant culture. Thank you for purchasing an authorised edition of this book and for respecting intellectual property laws by not reproducing, scanning or distributing any part of it by any means without permission. You are supporting authors and enabling Penguin Random House to continue to publish books for everyone. No part of this book may be used or reproduced in any manner for the purpose of training artificial intelligence technologies or systems. In accordance with Article 4(3) of the DSM Directive 2019/790, Penguin Random House expressly reserves this work from the text and data mining exception.

Printed and bound in Great Britain by Clays Ltd, Elcograf S.p.A.

The authorised representative in the EEA is Penguin Random House Ireland,
Morrison Chambers, 32 Nassau Street, Dublin D02 YH68

A CIP catalogue record for this book is available from the British Library

ISBN 9781529933611

Penguin Random House is committed to a sustainable future
for our business, our readers and our planet. This book is made
from Forest Stewardship Council® certified paper.



Copyright Material

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

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

Copyrighted Material

Copyrighted Material

Contents

PROLOGUE	xi
----------	----

PART I: Human Networks

CHAPTER 1: What Is Information?	3
CHAPTER 2: Stories: Unlimited Connections	18
CHAPTER 3: Documents: The Bite of the Paper Tigers	40
CHAPTER 4: Errors: The Fantasy of Infallibility	70
CHAPTER 5: Decisions: A Brief History of Democracy and Totalitarianism	118

PART II: The Inorganic Network

CHAPTER 6: The New Members: How Computers Are Different from Printing Presses	193
CHAPTER 7: Relentless: The Network Is Always On	231
CHAPTER 8: Fallible: The Network Is Often Wrong	257

PART III: Computer Politics

CHAPTER 9: Democracies: Can We Still Hold a Conversation?	305
CHAPTER 10: Totalitarianism: All Power to the Algorithms?	348
CHAPTER 11: The Silicon Curtain: Global Empire or Global Split?	361
EPILOGUE	395
AFTERWORD	405
ACKNOWLEDGEMENTS	417
NOTES	421
INDEX	491

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 an answer 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 popularised 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 axe, 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, and when people are fed bad information, they are likely to make bad decisions, no matter how wise and kind they personally are. 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 the Second World War. Stalin eventually emerged as the victor of that war,¹ 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 analyse 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 judge-

ments 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.² 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.'³

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.'⁴ 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'.⁵

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'.⁶

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.⁷

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 per cent, but in the 1790s it had an average mortality rate of 50 per cent, and for Goethe's four younger children it was a death sentence.⁸

Altogether in Goethe's family – a well-to-do German family in the late eighteenth century – the child survival rate was an abysmal 25 per cent. 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 per cent of German children reached age fifteen,⁹ and the same was probably true in most other parts of the world.¹⁰ By 2020, 95.6 per cent of children worldwide lived beyond their fifteenth birthday,¹¹ and in Germany that figure was 99.5 per cent.¹² This momentous achievement would not have been possible without collecting, analysing 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 Alexandria¹³ 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 jeopardise 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 6 June 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.'¹⁴

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 analyses them at length, but believes they could be mitigated successfully.¹⁵

Others are more sceptical. 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 civilisation.¹⁶ 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'.¹⁷ In a 2023 survey of 2,778 AI researchers, more than a third gave at least a 10 per cent chance of advanced AI leading to outcomes as bad as human extinction.¹⁸ In 2023 close to thirty governments – including those of China, the United States and the UK – 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'.¹⁹ 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 re-engineer 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 synthesise 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.

WEAPONISING 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 hypothesised 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. As AI evolves and accumulates power, the human-centred system we have created will be replaced by a radically different system that treats all phenomena, including even human beings, as clusters and eddies of information. In the human-centred system, authority stems from human feelings and desires, and the highest values are the rights and freedoms of human beings. *Homo Deus*

warned that in a system that understands everything in the universe as information flows, human feelings will be a target for hacking, human desires will be a target for manipulation, and the rights of human beings will be trumped in the name of a radically new ideal – the freedom of information. Whereas freedom of speech has been a traditional right granted to human beings, freedom of information is something granted to informational currents and to non-human informational entities like AIs.

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. Human feelings and desires are hacked and manipulated like never before, whereas information moves and evolves so fast that it threatens to make the world unintelligible and unmanageable for human cognition. 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 anti-establishment 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.²⁰

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 world view, 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 sceptical 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.²¹ 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 recognise and correct their own mistakes.²²

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’.²³ 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”’.²⁴ 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 anti-establishment movements in previous genera-

tions. 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 sceptical 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.²⁵ 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 organisation. 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 characterised 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.²⁶

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 God²⁷ 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.²⁸ 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 in 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 destabilise 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 centralised 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 USSR 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 canonised provides valuable insight about present-day claims for AI infallibility. Similarly, studying the early modern witch hunts and Stalin's collectivisation offers stark warnings about what might go wrong as we give AIs greater control over twenty-first-century societies. A deep knowledge of his-

tory 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 emphasised 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 emphasises the shift from organic to inorganic information networks. The Roman Empire, the Catholic Church and the USSR 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 terrorising and censoring their own agents. But how can a human dictator terrorise 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 defence 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?

PART I

Human Networks

Copyrighted Material

Copyrighted Material

CHAPTER 1

What Is Information?

It is always tricky to define fundamental concepts. Since they are the basis for everything that follows, they themselves seem to lack any basis of their own. Physicists have a hard time defining matter and energy, biologists have a hard time defining life, and philosophers have a hard time defining reality.

Information is increasingly seen by many philosophers and biologists, and even by some physicists, as the most basic building block of reality, more elementary than matter and energy.¹ No wonder that there are many disputes about how to define information, and how it is related to the evolution of life or to basic ideas in physics such as entropy, the laws of thermodynamics and the quantum uncertainty principle.² This book will make no attempt to resolve – or even explain – these disputes, nor will it offer a universal definition of information applicable to physics, biology and all other fields of knowledge. Since it is a work of history, which studies the past and future development of human societies, it will focus on the definition and role of information in history.

In everyday usage, information is associated with human-made symbols like spoken or written words. Consider, for example, the

story of Cher Ami and the Lost Battalion. In October 1918, when the American Expeditionary Forces were fighting to liberate northern France from the Germans, a battalion of more than five hundred American soldiers was trapped behind enemy lines. American artillery, which was trying to provide them with cover fire, misidentified their location and dropped the barrage directly on them. The battalion's commander, Major Charles Whittlesey, urgently needed to inform headquarters of his true location, but no runner could break through the German line. According to several accounts, as a last resort Whittlesey turned to Cher Ami, an army carrier pigeon. On a tiny piece of paper, Whittlesey wrote, 'We are along the road parallel [*sic*] 276.4. Our artillery is dropping a barrage directly on us. For heaven's sake stop it.' The paper was inserted into a canister on Cher Ami's right leg, and the bird was released into the air. One of the battalion's soldiers, Private John Nell, recalled years later, 'We knew without a doubt this was our last chance. If that one lonely, scared pigeon failed to find its loft, our fate was sealed.'

Witnesses later described how Cher Ami flew into heavy German fire. A shell exploded directly below the bird, killing five men and severely injuring the pigeon. A splinter tore through Cher Ami's chest, and his right leg was left hanging by a tendon. But he got through. The wounded pigeon flew the forty kilometres to division headquarters in about forty-five minutes, with the canister containing the crucial message attached to the remnant of his right leg. Though there is some controversy about the exact details, it is clear that the American artillery adjusted its barrage, and an American counterattack rescued the Lost Battalion. Cher Ami was tended by army medics, sent to the United States as a hero, and became the subject of numerous articles, short stories, children's books, poems and even movies. The pigeon had no idea what information he was conveying, but the symbols inked on the piece of paper he carried helped save hundreds of men from death and captivity.³

Information, however, does not have to consist of human-made symbols. According to the biblical myth of the Flood, Noah learned

that the water had finally receded because the pigeon he sent out from the ark returned with an olive branch in her mouth. Then God set a rainbow in the clouds as a heavenly record of his promise never to flood the earth again. Pigeons, olive branches and rainbows have since become iconic symbols of peace and tolerance. Objects that are even more remote than rainbows can also be information. For astronomers the shape and movement of galaxies constitute crucial information about the history of the universe. For navigators the North Star indicates which way is north. For astrologers the stars are a cosmic script, conveying information about the future of individual humans and entire societies.

Of course, defining something as ‘information’ is a matter of perspective. An astronomer or astrologer might view the Libra constellation as ‘information’, but these distant stars are far more than just a notice board for human observers. There might be an alien civilisation up there, totally oblivious to the information we glean from their home and to the stories we tell about it. Similarly, a piece of paper marked with ink splotches can be crucial information for an army unit, or dinner for a family of termites. Any object can be information – or not. This makes it difficult to define what information is.

The ambivalence of information has played an important role in the annals of military espionage, when spies needed to communicate information surreptitiously. During the First World War, northern France was not the only major battleground. From 1915 to 1918 the British and Ottoman Empires fought for control of the Middle East. After repulsing an Ottoman attack on the Sinai Peninsula and the Suez Canal, the British in turn invaded the Ottoman Empire, but were held at bay until October 1917 by a fortified Ottoman line stretching from Beersheba to Gaza. British attempts to break through were repulsed at the First Battle of Gaza (26 March 1917) and the Second Battle of Gaza (17–19 April 1917). Meanwhile, pro-British Jews living in Palestine set up a spy network code-named NILI to inform the British about Ottoman troop movements. One method they developed to communicate with their British operators involved

window shutters. Sarah Aaronsohn, a NILI commander, had a house overlooking the Mediterranean. She signalled British ships by closing or opening a particular shutter, according to a predetermined code. Numerous people, including Ottoman soldiers, could obviously see the shutter, but nobody other than NILI agents and their British operators understood it was vital military information.⁴ So, when is a shutter just a shutter, and when is it information?

The Ottomans eventually caught the NILI spy ring due in part to a strange mishap. In addition to shutters, NILI used carrier pigeons to convey coded messages. On 3 September 1917, one of the pigeons diverged off course and landed in – of all places – the house of an Ottoman officer. The officer found the coded message but couldn't decipher it. Nevertheless, the pigeon itself was crucial information. Its existence indicated to the Ottomans that a spy ring was operating under their noses. As Marshall McLuhan might have put it, the pigeon was the message. NILI agents learned about the capture of the pigeon and immediately killed and buried all the remaining birds they had, because the mere possession of carrier pigeons was now incriminating information. But the massacre of the pigeons did not save NILI. Within a month the spy network was uncovered, several of its members were executed, and Sarah Aaronsohn committed suicide to avoid divulging NILI's secrets under torture.⁵ When is a pigeon just a pigeon, and when is it information?

Clearly, then, information cannot be defined as specific types of material objects. Any object – a star, a shutter, a pigeon – can be information in the right context. So exactly what context defines such objects as 'information'? The naive view of information argues that objects are defined as information in the context of truth seeking. Something is information if people use it to try to discover the truth. This view links the concept of information with the concept of truth and assumes that the main role of information is to represent reality. There is a reality 'out there', and information is something that represents that reality and that we can therefore use to learn about real-

ity. For example, the information NILI provided the British was meant to represent the reality of Ottoman troop movements. If the Ottomans massed ten thousand soldiers in Gaza – the centrepiece of their defences – a piece of paper with symbols representing ‘ten thousand’ and ‘Gaza’ was important information that could help the British win the battle. If, on the other hand, there were actually twenty thousand Ottoman troops in Gaza, that piece of paper did not represent reality accurately, and could lead the British to make a disastrous military mistake.

Put another way, the naive view argues that information is an attempt to represent reality, and when this attempt succeeds, we call it truth. While this book takes many issues with the naive view, it agrees that truth is an accurate representation of reality. But this book also holds that most information is *not* an attempt to represent reality and that what defines information is something entirely different. Most information in human society, and indeed in other biological and physical systems, *does not represent anything*.

I want to spend a little longer on this complex and crucial argument, because it constitutes the theoretical basis of the book.

WHAT IS TRUTH?

Throughout this book, ‘truth’ is understood as something that accurately represents certain aspects of reality. Underlying the notion of truth is the premise that there exists one universal reality. Anything that has ever existed or will ever exist in the universe – from the North Star, to the NILI pigeon, to web pages on astrology – is part of this single reality. This is why the search for truth is a universal project. While different people, nations or cultures may have competing beliefs and feelings, they cannot possess contradictory truths, because they all share a universal reality. Anyone who rejects universalism rejects truth.

Truth and reality are nevertheless different things, because no

matter how truthful an account is, it can never represent reality in all its aspects. If a NILI agent wrote that there are ten thousand Ottoman soldiers in Gaza, and there were indeed ten thousand soldiers there, this accurately pointed to a certain aspect of reality, but it neglected many other aspects. The very act of counting entities – whether apples, oranges, or soldiers – necessarily focuses attention on the similarities between these entities while discounting differences.⁶ For example, saying only that there were ten thousand Ottoman soldiers in Gaza neglected to specify whether some were experienced veterans and others were green recruits. If there were a thousand recruits and nine thousand old hands, the military reality was quite different from if there were nine thousand rookies and a thousand battle-hardened veterans.

There were many other differences between the soldiers. Some were healthy; others were sick. Some Ottoman troops were ethnically Turkish, while others were Arabs, Kurds, or Jews. Some were brave, others cowardly. Indeed, each soldier was a unique human being, with different parents and friends and individual fears and hopes. First World War poets like Wilfred Owen famously attempted to represent these latter aspects of military reality, which mere statistics never conveyed accurately. Does this imply that writing ‘ten thousand soldiers’ is always a misrepresentation of reality, and that to describe the military situation around Gaza in 1917, we must specify the unique history and personality of every soldier?

Another problem with any attempt to represent reality is that reality contains many viewpoints. For example, present-day Israelis, Palestinians, Turks and Britons have different perspectives on the British invasion of the Ottoman Empire, the NILI underground and the activities of Sarah Aaronsohn. That does not mean, of course, that there are several entirely separate realities, or that there are no historical facts. There is just one reality, but it is complex.

Reality includes an objective level with objective facts that don't depend on people's beliefs; for example, it is an objective fact that Sarah Aaronsohn died on 9 October 1917, from self-inflicted gun-

shot wounds. Saying that ‘Sarah Aaronsohn died in an airplane crash on 15 May 1919’ is an error.

Reality also includes a subjective level with subjective facts like the beliefs and feelings of various people, but in this case, too, facts can be separated from errors. For example, it is a fact that Israelis tend to regard Aaronsohn as a patriotic hero. Three weeks after her suicide, the information NILI supplied helped the British finally break the Ottoman line at the Battle of Beersheba (31 October 1917) and the Third Battle of Gaza (1–2 November 1917). On 2 November 1917, the British foreign secretary, Arthur Balfour, issued the Balfour Declaration, announcing that the British government ‘view with favour the establishment in Palestine of a national home for the Jewish people’. Israelis credit this in part to NILI and Sarah Aaronsohn, whom they admire for her sacrifice. It is another fact that Palestinians evaluate things very differently. Rather than admiring Aaronsohn, they regard her – if they’ve heard about her at all – as an imperialist agent. Even though we are dealing here with subjective views and feelings, we can still distinguish truth from falsehood. For views and feelings – just like stars and pigeons – are a part of the universal reality. Saying that ‘Sarah Aaronsohn is admired by everyone for her role in defeating the Ottoman Empire’ is an error, not in line with reality.

Nationality is not the only thing that affects people’s viewpoint. Israeli men and Israeli women may see Aaronsohn differently, and so do left-wingers and right-wingers, or Orthodox and secular Jews. Since suicide is forbidden by Jewish religious law, Orthodox Jews have difficulty seeing Aaronsohn’s suicide as a heroic act (she was actually denied burial in the hallowed ground of a Jewish cemetery). Ultimately, each individual has a different perspective on the world, shaped by the intersection of different personalities and life histories. Does this imply that when we wish to describe reality, we must always list all the different viewpoints it contains and that a truthful biography of Sarah Aaronsohn, for example, must specify how every single Israeli and Palestinian has felt about her?

Taken to extremes, such a pursuit of accuracy may lead us to try to represent the world on a one-to-one scale, as in the famous Jorge Luis Borges story 'On Exactitude in Science' (1946). In this story Borges tells of a fictitious ancient empire that became obsessed with producing ever more accurate maps of its territory, until eventually it produced a map with a one-to-one scale. The entire empire was covered with a map of the empire. So many resources were wasted on this ambitious representational project that the empire collapsed. Then the map too began to disintegrate, and Borges tells us that only 'in the western Deserts, tattered fragments of the map are still to be found, sheltering an occasional beast or beggar'.⁷ A one-to-one map may look like the ultimate representation of reality, but tellingly it is no longer a representation at all; it is the reality.

The point is that even the most truthful accounts of reality can never represent it in full. There are always some aspects of reality that are neglected or distorted in every representation. Truth, then, isn't a one-to-one representation of reality. Rather, truth is something that brings our attention to certain aspects of reality while inevitably ignoring other aspects. No account of reality is 100 per cent accurate, but some accounts are nevertheless more truthful than others.

WHAT INFORMATION DOES

As noted above, the naive view sees information as an attempt to represent reality. It is aware that some information doesn't represent reality well, but it dismisses this as unfortunate cases of 'misinformation' or 'disinformation'. Misinformation is an honest mistake, occurring when someone tries to represent reality but gets it wrong. Disinformation is a deliberate lie, occurring when someone consciously intends to distort our view of reality.

The naive view further believes that the solution to the problems caused by misinformation and disinformation is more information. This idea, sometimes called the counterspeech doctrine, is associated

with the US Supreme Court justice Louis D. Brandeis, who wrote in *Whitney v. California* (1927) that the remedy to false speech is more speech and that in the long term free discussion is bound to expose falsehoods and fallacies. If all information is an attempt to represent reality, then as the amount of information in the world grows, we can expect the flood of information to expose the occasional lies and errors and to ultimately provide us with a more truthful understanding of the world.

On this crucial point, this book strongly disagrees with the naive view. There certainly are instances of information that attempt to represent reality and succeed in doing so, but this is *not* the defining characteristic of information. A few pages ago I referred to stars as information and casually mentioned astrologers alongside astronomers. Adherents of the naive view of information probably squirmed in their chairs when they read it. According to the naive view, astronomers derive 'real information' from the stars, while the information that astrologers imagine to read in constellations is either 'misinformation' or 'disinformation'. If only people were given more information about the universe, surely they would abandon astrology altogether. But the fact is that for thousands of years astrology has had a huge impact on history, and today millions of people still check their star signs before making the most important decisions of their lives, like what to study and whom to marry. As of 2021, the global astrology market was valued at \$12.8 billion.⁸

No matter what we think about the accuracy of astrological information, we should acknowledge its important role in history. It has connected lovers, and even entire empires. Roman emperors routinely consulted astrologers before making decisions. Indeed, astrology was held in such high esteem that casting the horoscope of a reigning emperor was a capital offence. Presumably, anyone casting such a horoscope could foretell when and how the emperor would die.⁹ Rulers in some countries still take astrology very seriously. In 2005 the junta of Myanmar allegedly moved the country's capital from Yangon to Naypyidaw based on astrological advice.¹⁰ A theory

of information that cannot account for the historical significance of astrology is clearly inadequate.

What the example of astrology illustrates is that errors, lies, fantasies and fictions are information, too. Contrary to what the naive view of information says, information has no essential link to truth, and its role in history isn't to represent a pre-existing reality. Rather, what information does is to create *new* realities by tying together disparate things – whether couples or empires. Its defining feature is connection rather than representation, and information is whatever connects different points into a network. Information doesn't necessarily inform us about things. Rather, it puts things in formation. Horoscopes put lovers in astrological formations, propaganda broadcasts put voters in political formations and marching songs put soldiers in military formations.

As a paradigmatic case, consider music. Most symphonies, melodies and tunes don't represent anything, which is why it makes no sense to ask whether they are true or false. Over the years people have created a lot of bad music, but not fake music. Without representing anything, music nevertheless does a remarkable job in connecting large numbers of people and synchronising their emotions and movements. Music can make soldiers march in formation, clubbers sway together, church congregations clap in rhythm and sports fans chant in unison.¹¹

The role of information in connecting things is of course not unique to human history. A case can be made that this is the chief role of information in biology, too.¹² Consider DNA, the molecular information that makes organic life possible. Like music, DNA doesn't represent reality. Though generations of zebras have been fleeing lions, you cannot find in the zebra DNA a string of nucleobases representing 'lion' nor another string representing 'flight'. Similarly, zebra DNA contains no representation of the sun, wind, rain or any other external phenomena that zebras encounter during their lives. Nor does DNA represent internal phenomena like body

organs or emotions. There is no combination of nucleobases that represents a heart or fear.

Instead of trying to represent pre-existing things, DNA helps to produce entirely new things. For instance, various strings of DNA nucleobases initiate cellular chemical processes that result in the production of adrenaline. Adrenaline too doesn't represent reality in any way. Rather, adrenaline circulates through the body, initiating additional chemical processes that increase the heart rate and direct more blood to the muscles.¹³ DNA and adrenaline thereby help to connect trillions of cells in the heart, legs and other body parts to form a functioning network that can do remarkable things, like run away from a lion.

If DNA represented reality, we could have asked questions like 'Does zebra DNA represent reality more accurately than lion DNA?' or 'Is the DNA of one zebra telling the truth, while another zebra is misled by her fake DNA?' These, of course, are nonsensical questions. We might evaluate DNA by the fitness of the organism it produces, but not by truthfulness. While it is common to talk about DNA 'errors', this refers only to mutations in the process of copying DNA – not to a failure to represent reality accurately. A mutation that inhibits the production of adrenaline reduces fitness, causing the network of cells to disintegrate, as when the zebra is killed and its trillions of cells lose connection with one another. But this kind of network failure means disintegration not disinformation. That's as true of countries, political parties and news networks as it is of zebras. Their existence too is jeopardised by loss of contact between their constituent parts, more than by inaccurate representations of reality.

Crucially, errors in the copying of DNA don't always reduce fitness. Once in a blue moon, they increase fitness. Without such mutations, there would be no process of evolution. All life-forms exist thanks to genetic 'errors'. The wonders of evolution are possible because DNA doesn't represent any pre-existing realities; it creates new realities.

Let us pause to digest the implications of this. Information is something that creates new realities by connecting different points into a network. This still includes the view of information as representation. Sometimes, a truthful representation of reality can connect humans, as when 600 million people sat glued to their television sets in July 1969, watching Neil Armstrong and Buzz Aldrin walking on the moon.¹⁴ The images on the screens accurately represented what was happening 384,000 kilometres away, and seeing them gave rise to feelings of awe, pride and human brotherliness that helped connect people.

However, such fraternal feelings can be produced in other ways, too. The emphasis on connection leaves ample room for other types of information that do not represent reality well. Sometimes erroneous representations of reality might also serve as a social nexus, as when millions of followers of a conspiracy theory watch a YouTube video claiming that the moon landing never happened. These images convey an erroneous representation of reality, but they might nevertheless give rise to feelings of anger against the establishment or pride in one's own wisdom that help create a cohesive new group.

Sometimes networks can be connected without *any* attempt to represent reality, neither accurate nor erroneous, as when genetic information connects trillions of cells or when a stirring musical piece connects thousands of humans.

As a final example, consider Mark Zuckerberg's vision of the Metaverse. The Metaverse is a virtual universe made entirely of information. Unlike the one-to-one map built by Jorge Luis Borges's imaginary empire, the Metaverse isn't an attempt to represent our world, but rather an attempt to augment or even replace our world. It doesn't offer us a digital replica of Buenos Aires or Salt Lake City; it invites people to build new virtual communities with novel landscapes and rules. As of 2024 the Metaverse seems like an overblown pipe dream, but within a couple of decades billions of people might migrate to live much of their lives in an augmented virtual reality, holding there most of their social and professional activities. People

might come to build relationships, join movements, hold jobs and experience emotional ups and downs in environments made of bits rather than atoms. Perhaps only in some remote deserts, tattered fragments of the old reality could still be found, sheltering an occasional beast or beggar.

INFORMATION IN HUMAN HISTORY

Viewing information as a social nexus helps us understand many aspects of human history that confound the naive view of information as representation. It explains the historical success not only of astrology but of much more important things, like the Bible. While some may dismiss astrology as a quaint sideshow in human history, nobody can deny the central role the Bible has played. If the main job of information had been to represent reality accurately, it would have been hard to explain why the Bible became one of the most influential texts in history.

The Bible makes many serious errors in its description of both human affairs and natural processes. The book of Genesis claims that all human groups – including, for example, the San people of the Kalahari Desert and the Aborigines of Australia – descend from a single family that lived in the Middle East about four thousand years ago.¹⁵ According to Genesis, after the Flood all Noah's descendants lived together in Mesopotamia, but following the destruction of the Tower of Babel they spread to the four corners of the Earth and became the ancestors of all living humans. In fact, the ancestors of the San people lived in Africa for hundreds of thousands of years without ever leaving the continent, and the ancestors of the Aborigines settled Australia more than fifty thousand years ago.¹⁶ Both genetic and archaeological evidence rule out the idea that the entire ancient populations of South Africa and Australia were annihilated about four thousand years ago by a flood and that these areas were subsequently repopulated by Middle Eastern immigrants.

An even graver distortion involves our understanding of infectious diseases. The Bible routinely depicts epidemics as divine punishment for human sins¹⁷ and claims they can be stopped or prevented by prayers and religious rituals.¹⁸ However, epidemics are of course caused by pathogens and can be stopped or prevented by following hygiene rules and using medicines and vaccines. This is today widely accepted even by religious leaders like the pope, who during the COVID-19 pandemic advised people to self-isolate, instead of congregating to pray together.¹⁹

Yet while the Bible has done a poor job in representing the reality of human origins, migrations and epidemics, it has nevertheless been very effective in connecting billions of people and creating the Jewish and Christian religions. Like DNA initiating chemical processes that bind billions of cells into organic networks, the Bible initiated social processes that bonded billions of people into religious networks. And just as a network of cells can do things that single cells cannot, so a religious network can do things that individual humans cannot, like building temples, maintaining legal systems, celebrating holidays and waging holy wars.

To conclude, information sometimes represents reality, and sometimes doesn't. But it always connects. This is its fundamental characteristic. Therefore, when examining the role of information in history, although it sometimes makes sense to ask 'How well does it represent reality? Is it true or false?' often the more crucial questions are 'How well does it connect people? What new network does it create?'

It should be emphasised that rejecting the naive view of information as representation does not force us to reject the notion of truth, nor does it force us to embrace the populist view of information as a weapon. While information always connects, some types of information – from scientific books to political speeches – may strive to connect people by accurately representing certain aspects of reality. But this requires a special effort, which most information does not make. This is why the naive view is wrong to believe that

creating more powerful information technology will necessarily result in a more truthful understanding of the world. If no additional steps are taken to tilt the balance in favour of truth, an increase in the amount and speed of information is likely to swamp the relatively rare and expensive truthful accounts with much more common and cheap types of information.

When we look at the history of information from the Stone Age to the Silicon Age, we therefore see a constant rise in connectivity, without a concomitant rise in truthfulness or wisdom. Contrary to what the naive view believes, *Homo sapiens* didn't conquer the world because we are talented at turning information into an accurate map of reality. Rather, the secret of our success is that we are talented at using information to connect lots of individuals. Unfortunately, this ability often goes hand in hand with believing in lies, errors and fantasies. This is why even technologically advanced societies like Nazi Germany and the Soviet Union have been prone to hold delusional ideas, without their delusions necessarily weakening them. Indeed, the mass delusions of Nazi and Stalinist ideologies about things like race and class actually helped them make tens of millions of people march together in lockstep.

In chapters 2–5 we'll take a closer look at the history of information networks. We'll discuss how, over tens of thousands of years, humans invented various information technologies that greatly improved connectivity and cooperation without necessarily resulting in a more truthful representation of the world. These information technologies – invented centuries and millennia ago – still shape our world even in the era of the internet and AI. The first information technology we'll examine, which is also the first information technology developed by humans, is the story.

CHAPTER 2

Stories: Unlimited Connections

We Sapiens rule the world not because we are so wise but because we are the only animals that can cooperate flexibly in large numbers. I have explored this idea in my previous books *Sapiens* and *Homo Deus*, but a brief recap is inescapable.

The Sapiens' ability to cooperate flexibly in large numbers has precursors among other animals. Some social mammals like chimpanzees display significant flexibility in the way they cooperate, while some social insects like ants cooperate in very large numbers. But neither chimps nor ants establish empires, religions or trade networks. Sapiens are capable of doing such things because we are far more flexible than chimps and can simultaneously cooperate in even larger numbers than ants. In fact, there is no upper limit to the number of Sapiens who can cooperate with one another. The Catholic Church has about 1.4 billion members. China has a population of about 1.4 billion. The global trade network connects about 8 billion Sapiens.

This is surprising given that humans cannot form long-term intimate bonds with more than a few hundred individuals.¹ It takes many years and common experiences to get to know someone's

unique character and history and to cultivate ties of mutual trust and affection. Consequently, if Sapiens networks were connected only by personal human-to-human bonds, our networks would have remained very small. This is the situation among our chimpanzee cousins, for example. Their typical community numbers 20–60 members, and on rare occasions the number might increase to 150–200.² This appears to have been the situation also among ancient human species like Neanderthals and archaic *Homo sapiens*. Each of their bands numbered a few dozen individuals, and different bands rarely cooperated.³

About seventy thousand years ago, *Homo sapiens* bands began displaying an unprecedented capacity to cooperate with one another, as evidenced by the emergence of inter-band trade and artistic traditions and by the rapid spread of our species from our African homeland to the entire globe. What enabled different bands to cooperate is that evolutionary changes in brain structure and linguistic abilities apparently gave Sapiens the aptitude to tell and believe fictional stories and to be deeply moved by them. Instead of building a network from human-to-human chains alone – as the Neanderthals, for example, did – stories provided *Homo sapiens* with a new type of chain: human-to-story chains. In order to cooperate, Sapiens no longer had to know each other personally; they just had to know the same story. And the same story can be familiar to billions of individuals. A story can thereby serve like a central connector, with an unlimited number of outlets into which an unlimited number of people can plug. For example, the 1.4 billion members of the Catholic Church are connected by the Bible and other key Christian stories; the 1.4 billion citizens of China are connected by the stories of communist ideology and Chinese nationalism; and the 8 billion members of the global trade network are connected by stories about currencies, corporations and brands.

Even charismatic leaders who have millions of followers are an example of this rule rather than an exception. It may seem that in the case of ancient Chinese emperors, medieval Catholic popes or

modern corporate titans it has been a single flesh-and-blood human – rather than a story – that has served as a nexus linking millions of followers. But, of course, in all these cases almost none of the followers has had a personal bond with the leader. Instead, what they have connected to has been a carefully crafted *story* about the leader, and it is in this story that they have put their faith.

Joseph Stalin, who stood at the nexus of one of the biggest personality cults in history, understood this well. When his troublesome son Vasily exploited his famous name to frighten and awe people, Stalin berated him. ‘But I’m a Stalin too,’ protested Vasily. ‘No, you’re not,’ replied Stalin. ‘You’re not Stalin and I’m not Stalin. Stalin is Soviet power. Stalin is what he is in the newspapers and the portraits, not you, no – not even me!’⁴

Present-day influencers and celebrities would concur. Some have hundreds of millions of online followers, with whom they communicate daily through social media. But there is very little authentic personal connection there. The social media accounts are usually run by a team of experts, and every image and word is professionally crafted and curated to manufacture what is nowadays called a brand.⁵

A ‘brand’ is a specific type of story. To brand a product means to tell a story about that product, which may have little to do with the product’s actual qualities but which consumers nevertheless learn to associate with the product. For example, over the decades the Coca-Cola corporation has invested tens of billions of dollars in advertisements that tell and retell the story of the Coca-Cola drink.⁶ People have seen and heard the story so often that many have come to associate a certain concoction of flavored water with fun, happiness and youth (as opposed to tooth decay, obesity and plastic waste). That’s branding.⁷

As Stalin knew, it is possible to brand not only products but also individuals. A corrupt billionaire can be branded as the champion of the poor; a bungling imbecile can be branded as an infallible genius; and a guru who sexually abuses his followers can be branded as a

chaste saint. People think they connect to the person, but in fact they connect to the story told *about* the person, and there is often a huge gulf between the two.

Even the story of Cher Ami, the heroic pigeon, was partly the product of a branding campaign aimed at enhancing the public image of the US Army's Pigeon Service. A 2021 revisionist study by the historian Frank Blazich found that though there is no doubt Cher Ami sustained severe injuries while transporting a message somewhere in Northern France, several key features of the story are doubtful or inaccurate. First, relying on contemporary military records, Blazich demonstrated that headquarters learned about the exact location of the Lost Battalion about twenty minutes *prior* to the pigeon's arrival. It was not the pigeon that put a stop to the barrage of friendly fire decimating the Lost Battalion. Even more crucially, there is simply no proof that the pigeon carrying Major Whittlesey's message was Cher Ami. It might well have been another bird, while Cher Ami might have sustained his wounds a couple of weeks later, during an altogether different battle.

According to Blazich, the doubts and inconsistencies in Cher Ami's story were overshadowed by its propaganda value to the army and its appeal to the public. Over the years the story was retold so many times that facts became hopelessly enmeshed with fiction. Journalists, poets and filmmakers added fanciful details to it, for example that the pigeon lost an eye as well as a leg and that it was awarded the Distinguished Service Cross. In the 1920s and 1930s Cher Ami became the most famous bird in the world. When he died, his carefully preserved corpse was placed on display at the Smithsonian's National Museum of American History, where it became a pilgrimage site for American patriots and World War I veterans. As the story grew in the telling, it took over even the recollections of survivors of the Lost Battalion, who came to accept the popular narrative at face value. Blazich recounts the case of Sherman Eager, an officer in the Lost Battalion, who decades after the war brought his