The Coming Wave — 20-Minute Read Summary

Mustafa Suleyman, with Michael Bhaskar (2023)

Daily Book, Episode 36 — from Bill Gates's reading list for 2024, where Gates called it the book on artificial intelligence he recommends the most.



The Book in 60 Seconds

A wave of powerful technologies — above all artificial intelligence and synthetic biology — is building now and will reshape medicine, work, war, and politics within a decade. These tools promise extraordinary abundance, but they share four dangerous traits: they can be used for almost anything, they improve at compounding speed, they give small actors outsized power, and they increasingly act on their own. Mustafa Suleyman, a co-founder of DeepMind writing from inside the industry he is warning about, calls the resulting challenge "the containment problem": can humanity keep control of technologies that get cheaper, more powerful, and more widespread at the same time? History says containment usually fails, or arrives painfully late. His answer is a "narrow path" between catastrophe (harms we fail to prevent) and dystopia (prevention through total surveillance), walked by building layered guardrails — technical safety, audits, choke-point controls, corporate responsibility, modern regulation, and international agreements — now, while the wave is still gathering.

The Author and Why He Matters

Most books warning about artificial intelligence are written by journalists, academics, or critics standing outside the industry. *The Coming Wave* is written by one of the industry's central builders. Suleyman co-founded DeepMind, the London laboratory behind landmark systems including AlphaFold, which solved the protein-folding problem, and which became part of Google. He later co-founded Inflection AI and went on to lead consumer artificial intelligence at Microsoft. His co-author, Michael Bhaskar, is a writer and publisher who helps give the book its unusually readable, historically grounded prose.

That biography creates the book's defining tension. Suleyman is not warning about someone else's technology. He is warning about his own life's work — and continuing to do that work while he warns. Readers will differ on whether that makes him more credible (he knows what is actually coming, and has seen capabilities before the public does) or less (he and his employers stand to benefit from the rules he proposes). The book is most fairly read with both reactions in mind.

The Central Argument

Waves are how technology changes the world

Suleyman's framing device is the wave: a general-purpose technology that spreads through society in a surge, transforming everything it touches. The printing press, the steam engine, electricity, and the car were waves. Each brought large benefits, large harms, and long, messy periods in which societies slowly invented rules to manage them — licensing, safety standards, liability, and norms.

The coming wave, he argues, is bigger and faster than its predecessors because it is not one technology but two converging ones:

1. **Artificial intelligence** — systems that learn, generate, predict, and increasingly act, applicable to almost any task involving information.
2. **Synthetic biology** — the ability to read, write, and redesign living systems: editing genes with tools such as CRISPR, synthesising DNA at falling prices, and engineering organisms for medicine, agriculture, and industry.

The convergence is the point Bill Gates seized on in recommending the book. Artificial intelligence is accelerating biology: AlphaFold turned a decades-old scientific bottleneck into a solved problem almost overnight, and AI-designed proteins, drugs, and experiments are following. Biology, in return, gives artificial intelligence a physical frontier — laboratories, organisms, and eventually bodies — far beyond the screen.

The containment problem

"Containment" is Suleyman's name for humanity's ability to control, limit, and if necessary shut down a technology. The book's historical survey is sobering. Containment has worked best where a technology was **expensive, visible, and infrastructure-heavy**. Nuclear weapons are the standing example: they required rare materials, vast plants, and state-scale budgets, so treaties, inspections, and export controls could keep the club of nuclear states small.

The coming wave has the opposite profile. Its costs fall every year. Its tools shrink from industrial plants to laptops and benchtop devices. Its knowledge spreads through open publication and open-source release. Much of it is *dual-use by design*: the same model that summarises documents can write phishing campaigns; the same biological knowledge that designs a vaccine informs how a pathogen works. Suleyman's conclusion is that containment in the old sense — keeping the technology in a few trusted hands — is probably impossible. The question is whether a new kind of containment can be built instead.

The Key Ideas, Explained Plainly

1. Four traits make this wave hard to contain

Suleyman gives readers a diagnostic checklist. The coming wave is:

- **Omni-use.** It can be applied to almost any purpose, beneficial or harmful, by almost any user. There is no single "use case" to license or ban.
- **Hyper-evolutionary.** It improves fast, and improvement compounds: each generation of AI helps design the next, and each biological tool makes the next experiment cheaper. Regulation written for today's capability is outdated on arrival.
- **Asymmetric.** It lets small actors do what once required states or large organisations. Automated cyberattacks, cloned-voice fraud, and — in the biological case — potentially engineered pathogens all lower the team size needed to cause large harm.
- **Autonomous.** Systems increasingly act without step-by-step human supervision, at speeds and scales no human review process can match. Delegation is adopted one reasonable step at a time, until the cumulative delegation is something nobody explicitly chose.

Any one of these traits would strain governance. Together, they define the dilemma.

2. The pessimism-aversion trap

Human beings and institutions resist believing bad news about technologies they love, profit from, or depend on. Suleyman calls this reluctance the pessimism-aversion trap, and it explains a recurring pattern — with cars, fossil fuels, and social media, societies recognised harms late, after the technology was woven into daily life and reform had become expensive. His plea is to invert the sequence this once: build the guardrails during the design phase, not after the disaster phase.

3. The risks are concrete, not just cinematic

The book is sometimes misread as a treatise on far-future superintelligence. Its risk catalogue is mostly nearer and more mundane:

- **Cyberattacks** that are automated, personalised, and cheap to run at scale.
- **Disinformation and fraud** industrialised by generative systems — deepfake video, cloned voices, and synthetic text that erode shared trust in evidence.
- **Autonomous weapons** that lower the attacker's human cost of conflict and blur accountability.
- **Engineered pathogens and lab accidents**, arising from malice, carelessness, or sheer volume of experiments as biological capability diffuses.
- **Rapid cognitive-job displacement.** Previous automation waves spread factory by factory; a new model can be rolled out to millions of knowledge workers overnight, compressing the adjustment period societies usually rely on.
- **Power concentration and power diffusion at once**: the most advanced systems concentrate power in a few companies and states, while copies and weaker versions diffuse power to everyone else — a combination that destabilises existing orders from both directions.

4. The benefits are real, and they are why "just stop" fails

Suleyman is not a reluctant technologist. The same wave, contained, could deliver personalised medicines, rapid vaccine design, drought-resistant crops, pollution-eating microbes, abundant clean energy, and a general acceleration of science. That promise is why he rejects both a moratorium and nostalgic rejection: stopping the wave, even if it were possible, would forfeit cures and climate solutions. The argument is not *whether* to have these technologies, but *under what controls*.

5. The nation-state paradox

Powerful, diffuse tools weaken the nation-state — it can no longer monopolise the most consequential capabilities — yet only states, cooperating, can set and enforce rules at scale. States are simultaneously less able and more necessary. Much of the book's political argument follows from this paradox, including his attention to the US–China technology race, where mutual fear of falling behind is itself one of the wave's engines.

6. The narrow path and the ten steps

Suleyman frames the choice as a **narrow path between catastrophe and dystopia**. Fail to contain, and harms accumulate. Contain by total surveillance and licensing of everything, and you build a panopticon that crushes the openness the technology was supposed to serve.

His proposed route is layered containment, sketched in a ten-step plan whose spirit matters more than its exact enumeration:

- **Technical safety and audits** — test powerful systems before release, build in shutdown and monitoring capability, and treat safety as engineering, not marketing.
- **Choke points** — govern what is still concentrated and visible: the most advanced chips and the plants that fabricate them, the largest training runs and data centres, DNA-synthesis screening for dangerous sequences.
- **Maker responsibility** — the companies and researchers who build these systems bear explicit responsibility for foreseeable misuse, with business incentives that reward caution rather than pure speed.
- **Government that can keep up** — regulation designed for technologies that change monthly, with real technical expertise inside the state.
- **International agreements** — narrow, achievable bargains between rivals first (for example, norms against autonomous nuclear decision-making and against engineered pathogens), following the Cold War precedent that adversaries can cooperate on dangers they share.
- **Culture and public voice** — a professional culture that treats caution as maturity, and a public that gets a real say in the trade-offs, not just a choice between products.

He is candid that this programme is ambitious to the point of looking utopian, and that no single layer suffices. The claim is cumulative: many imperfect layers, together, might keep societies on the narrow path most of the time.

How the Book Unfolds

The book is built in three movements. The first establishes the pattern: a short history of technology waves and of the many attempts, from medieval crossbow bans to nuclear non-proliferation, to stop a powerful technology from spreading. The pattern Suleyman extracts is uncomfortable. Technologies that are useful, profitable, or militarily valuable almost always spread; the realistic choices are about *how* they spread, under what norms, and with what safeguards — not whether they spread at all. Anyone who begins the book expecting a proposal to uninvent artificial intelligence or gene editing will have that expectation corrected early.

The second movement profiles the wave itself. Separate chapters take artificial intelligence and synthetic biology in turn, explaining each in plain language for a general reader — what a large model actually is, why scale changed what these systems can do, what CRISPR and DNA synthesis made cheap, and why the two fields keep accelerating each other. Worked examples do much of the persuasion here: protein folding collapsing from a fifty-year problem to a database lookup; voice and video synthesis moving from studio trick to phone app; the price of reading and writing genetic code falling along a curve that veterans compare to the fall in computing costs. The tone is deliberately double — wonder and warning in alternating paragraphs — because Suleyman's whole argument depends on the reader holding both reactions at once.

The third movement turns to consequences and response. Chapters on labour, war, fraud, and state power work through the second-order effects — not just what the tools do, but what they do to the institutions built for a slower world. Only then does the book arrive at containment as a programme and at its closing appeal: that the 2020s are a brief window in which the largest systems are still built by a countable number of organisations, in a countable number of places, and that the window will not stay open.

Three Scenarios to Keep in Mind

Suleyman rarely uses the language of formal scenarios, but three futures structure the book, and naming them helps a reader follow it.

The muddle-through future is the historical default: harms occur, scandals force partial rules, industries adapt, and societies end up with a patchwork — some real safety gains, some theatre, and a permanent background level of fraud, accidents, and displacement that everyone learns to tolerate, much as we tolerate road deaths. The book's fear is that this future, acceptable for cars, may not be acceptable when the harms can be self-spreading (a pathogen) or system-corroding (industrialised disinformation).

The catastrophe future is not a single robot apocalypse. It is a stack of failures arriving together: an engineered or leaked pathogen, automated attacks on infrastructure, autonomous weapons lowering the threshold of a regional war, and information systems too polluted to coordinate a response. Suleyman's argument is that the *probability of each element rises with diffusion*, and that a politics already weakened by distrust would struggle to improvise containment mid-crisis.

The dystopia future is catastrophe's mirror image: safety achieved, at the price of pervasive monitoring of laboratories, computers, communications, and movement, with licensing regimes that decide who may think, publish, or experiment. He insists this is not a straw man — it is the path frightened publics reliably demand after a disaster, and the path some states are already equipped to walk. The "narrow path" is his name for refusing both: enough control to prevent the worst, enough openness to remain a free society, maintained indefinitely rather than solved once.

Concepts Worth Taking Away

- **General-purpose technology.** A technology, like electricity, whose main effect is to change thousands of other activities. General-purpose technologies cannot be governed use-by-use; they have to be governed at the level of inputs, standards, and liability.
- **Diffusion.** The spread of a capability from its inventors to everyone else. Suleyman treats diffusion as close to a law of technological history, powered by falling costs, published research, talent movement, and competition.
- **Choke point.** A place where a diffuse system still passes through a narrow gate — advanced chip fabrication, large compute clusters, DNA-synthesis providers. Choke points are where governance retains leverage.
- **Dual use.** Research or tools with both beneficial and harmful applications, where intent, not the artefact, determines the category. Dual use is why simple bans keep failing in biology and, increasingly, in AI.
- **The window.** The present period, in which frontier capability is concentrated enough to govern. Much of the book's urgency comes from the claim that the window closes as capability diffuses — rules written later govern a world that no longer exists.

What Has Happened Since 2023

Because the book was published in 2023, a 2026 reader can begin grading it — an exercise Suleyman's framework invites.

On the credit side: frontier models have continued to gain capability while the cost of using a given level of capability has fallen sharply; capable open-weight models have spread exactly along the diffusion curve the book describes; voice-cloning fraud and synthetic-media scams have moved from novelty to police-report routine; AI assistants have entered ordinary white-collar workflows at a speed that has made the labour chapters feel less speculative; and governments have begun, haltingly, to build the institutional furniture the book called for — AI safety institutes, frontier-model reporting requirements, and export controls on advanced chips, which are choke-point policy in everything but name.

On the debit side: the gravest biological scenarios remain unrealised, long-horizon fully autonomous agents have proved harder to make reliable than the most breathless 2023 commentary implied, and international cooperation has, if anything, thinned as technology competition has intensified. The book's directional claims have aged better than anyone's timeline claims usually do; its warnings should still be read as warnings, not as a record.

Perhaps the most interesting test is institutional. In 2023, "containment" was a metaphor in search of machinery. By 2026 there is at least the outline of machinery — evaluations before deployment, incident reporting, synthesis screening, compute thresholds — contested, incomplete, and uneven across countries, but real. Whether it thickens into something like aviation safety, or dissolves under competitive pressure, is the live question the book hands to its readers.

Questions for Discussion

1. Which of the four traits — omni-use, hyper-evolution, asymmetry, autonomy — do you find most convincing as a reason this wave differs from earlier ones? Which do you find weakest?
2. Does Suleyman's position inside the industry strengthen or weaken his argument for you? What standard should apply to warnings from builders generally?
3. If you could enact exactly one of his containment layers and no others — technical safety, choke points, maker liability, national regulation, or international agreement — which would you choose, and what would you expect it to prevent?
4. Where is your own line on the narrow path? Name a surveillance or licensing measure you would accept after a major AI or biological incident, and one you would refuse even then.
5. Which technology in your own working life is currently in its "pre-guardrails" phase, and what would its seat belt — a cheap, boring, standard safety practice — look like?

How to Approach Reading It

Read *The Coming Wave* as a **map, not a manual**. The first half — the wave framing, the history of containment, and the four traits — is the book's durable contribution, and it rewards slow reading. The second half is best read as a provocation to argue with: go through the ten steps and decide which you would actually support, which you would resist, and why.

It also helps to read the book **alongside the news**. For each story about a new model, a gene-editing breakthrough, an AI-powered scam, or a new technology law, ask the book's master question: does this make the wave more contained or less — and does it push us toward catastrophe, toward dystopia, or along the narrow path?

A final tip: notice your own resistance. Where you think "that will never happen," check whether that is evidence or optimism. Where you think "that is obviously inevitable," check for the opposite bias. The book is written to be tested, not believed.

Memorable Lines

Short, attributed excerpts only — the book rewards reading in full.

- Bill Gates, recommending the book: it is "the one I recommend the most" among books on AI.
- Suleyman's framing question, in paraphrase of his recurring formulation: can we forge "a narrow path between catastrophe and dystopia"?
- On containment, Suleyman argues it must become "a balance of power not between competing actors but between humans and our tools" — regulation, technical safety, governance, and accountability as one programme rather than separate fixes.

Fair Criticisms

The insider problem. Suleyman co-founded the field and now leads AI efforts at one of the world's largest technology companies. Critics reasonably ask whether containment, as he defines it, would entrench incumbents — who can afford audits, compliance, and licensing — against smaller competitors and open-source developers, and whether a builder's warning doubles as a builder's shield.

Utopian dependence on cooperation. Reviewers in publications such as *The Guardian* described the destination of the ten-step plan as effectively a utopian dream, however eloquently argued, in a world where the United States and China are competing rather than cooperating, and where treaties take years while models take months. Suleyman largely concedes the difficulty; the disagreement is over whether conceding it is enough.

A shifting definition. Some reviewers note that "containment" means something close to total control when the book recounts history, and something closer to ordinary regulation by the time the plan arrives. If containment means regulation, the argument is less revolutionary — societies do regulate cars, chemicals, guns, and medicines, imperfectly but often effectively — and the book's drama partly depends on the stronger early definition.

Having it both ways on optimism. The book expects AI to help solve climate change and disease, yet doubts it can be safely managed. Critics ask why a technology powerful enough for the first should not also be recruited for the second — and whether the optimism and the alarm are held to the same evidential standard.

Diagnosis over blueprint. The plan is a framework, not a manual. What exactly should a regulator do next week? How do you audit a model that updates continuously? Where exactly is the line between dangerous and legitimate biology? Readers wanting operational detail will finish the book persuaded of the problem and still improvising the solution.

Timing risk. Like all technology forecasting, some of Suleyman's timelines may prove too fast (the worst biological scenarios have not materialised) or too slow (capabilities and deepfake fraud arrived quickly). His defence, and it is a fair one, is that the direction — cheaper, faster, more diffuse, more autonomous — matters more than any single date.

Why It's on the List

Bill Gates placed *The Coming Wave* on his 2024 reading list and called it his favourite book on artificial intelligence — specifically because of the convergence insight. In Gates's summary, AI is "only one part of an unprecedented convergence of scientific breakthroughs," with gene editing and DNA synthesis racing forward in parallel, driven by falling costs, growing computing power, and competitive incentives among countries, companies, and individuals that make the advance "essentially unstoppable — and harder to control." For a reader whose philanthropy centres on disease, agriculture, and energy, Suleyman's insistence that the biology half matters as much as the software half is precisely the point. The book also became a *Financial Times* Business Book of the Year 2023 shortlistee and a book of the year in several major publications, and it is frequently cited in AI-governance debates as the insider account of the containment idea.

Read in 2026, it functions as a checkable forecast: much of its direction has held — cheaper and more capable models, spreading open systems, everyday synthetic-media fraud — while its worst cases remain, fortunately, warnings rather than history. That combination is a good reason to read it now rather than later.

A word on style, because it affects how the book should be weighed. This is a trade book written for decision-makers and general readers, not a technical monograph: there are no equations, the biology and machine learning are explained through analogy and example, and the historical chapters range widely and quickly. Specialists in either field will find simplifications to quarrel with, and the endnotes matter for anyone who wants to check a specific claim. But the simplification is also the service. The containment problem sits exactly at the junction of two technical fields, corporate strategy, and statecraft, and almost nobody's job description covers all four. A book that a hospital administrator, a chip executive, a legislator, and a graduate student can all read and then argue about in shared vocabulary is doing something genuinely scarce — which is one reason it travelled so far beyond the technology press, onto the *Financial Times* Business Book of the Year shortlist and into the reading lists of people, like Gates, who allocate money against the future it describes.

For the purposes of this commute series, it also pairs unusually well with two earlier episodes. Against Yuval Noah Harari's *Nexus* (Episode 21), which treats artificial intelligence primarily as an information and power problem, Suleyman adds the biological half and the builder's perspective. Against Klein and Thompson's *Abundance* (coming later in this queue), which asks why rich societies stopped building physical things, Suleyman asks what happens when the most powerful things become almost weightless — copyable at near-zero cost. Reading the three together gives you the optimistic, the civilisational, and the engineering views of the same decade.

What to Read Next

- **Superintelligence** by Nick Bostrom — the deeper, denser, more philosophical treatment of advanced AI risk that Suleyman's nearer-term account complements.
- **Abundance** by Ezra Klein and Derek Thompson — covered earlier in this series; the optimistic, build-more counterpart on what societies should do with new capability.
- **Nexus** by Yuval Noah Harari — covered in this series as Episode 21; a historian's take on information networks, AI, and power.
- **The Anxious Generation** by Jonathan Haidt — also on Gates's 2024 list, and next-but-one in this queue; what happened the last time a powerful technology reached children before society built its guardrails.

In one line: the wave is coming either way; the only real choice is whether we meet it with guardrails we designed, or with emergency powers we improvise.