

MUSTAFA SULEYMAN, *THE COMING WAVE: TECHNOLOGY, POWER, AND THE 21ST CENTURY'S GREATEST DILEMMA* (CROWN, 2023). 352 PP.

*Book Review by Rick Campanelli**

Introduction

Just a decade ago, physicist Stephen Hawking upset techno-optimists when he said the greatest existential threats to humanity come from science and technology. A year later—upsetting a whole different faction—Hawking said that only a world government could save the world from these threats. In *The Coming Wave*, Mustafa Suleyman provides a compelling explanation of how much and how quickly Artificial Intelligence (AI) ups the ante on the kinds of threats Hawking anticipated, arguing we need to find new systems that can adequately and quickly respond to constrain them. The book also begs a question: in responding to these new categories of risk, will we be more willing to embrace solutions that undermine democracy and human flourishing?

Suleyman is the consummate AI insider. He co-founded DeepMind and Inflection AI—the former became Google's AI powerhouse, and Microsoft is a primary if not dominant investor in the latter. In *The Coming Wave*, Suleyman offers an understandable explanation of just what Artificial General Intelligence (AGI) is, and how it creates vast, transformational potential on an unparalleled scale. His insider expertise makes the book clear and credible. Unfortunately, it also makes it difficult to dismiss its warnings. And Suleyman is worried. He warns that we are entering “an era when unprecedented opportunities . . . [are] matched by unprecedented risks.” Is he right? Reading *The Coming Wave* is a good way to get a handle on what AI is about, and what's at stake.

Overview

In the first part of the book, Suleyman describes how in just over a decade he came to appreciate the potential for AGI. He walks us through continuously accelerating computing leaps by teams at DeepMind, OpenAI, Google, Facebook, Microsoft, and the relatively small array of actors capable of amassing the vast resources needed to build these new systems, and that dominate the playing field. He explains how initial successes exponentially expanded AI capability and capacity; how large language models (LLMs—which he also explains) opened the flood gates; and how AI systems went from learning from human-curated inputs to creating and developing their own learning models and strategies. “A key ingredient of the LLM revolution,” he explains, “is that for the first time very large models could be trained directly, without the need for carefully curated and human-labeled data sets.” We watch the technologies transmute from fun gaming competitions into serious national security arms races among nations, which quickly realized they were on the brink of new kinds of wars. We are already seeing these waves beginning to break in Ukraine and Israel.

These chapters alone make the book worth reading, just to understand what we really are, or should be, talking about. But what also emerges is a portrait of highly enthusiastic and intelligent engineers, single-mindedly competing at break-neck speed to be on the ground floor, to lead and to be enriched by what they suspect may be the greatest technological wave in human history. And while Suleyman describes lots of conversations among tech leaders about practical and

* Lecturer in Public Policy, University of Virginia Batten School of Leadership & Public Policy; Director of Graduate Ministries, Center for Christian Study (Charlottesville). In the past, Rick served as Director of the Office for Civil Rights and Counselor to the Secretary, U.S. Department of Health & Human Services.

ethical implications, he worries there is no corresponding passion, incentive, or mechanism to appropriately address the accompanying risks—most of which will not be borne by the developers but by all of us. That question, perhaps the most important question, of whether and how to address the disruptive risks that accompany the massive potential benefits of AI is left to others—whoever they are.

What’s the Big Deal About AI?

We might have read articles on how AI touches on specific areas of interest to us, and its tremendous potential to disruptively research, write, strategize, and produce improvements ranging from unimagined therapies and diagnostics, to advances in agricultural and national defense capabilities. If we have, we also might have a lurking sense that each of these varied threads fails to reflect the big picture of the AI tapestry. According to Suleyman, we should pay attention to that sense.

We might wonder, for instance, will lawyers, doctors, farmers, researchers, coders, store clerks, and factory workers still have jobs? Will students be able to cheat undetectably on tests, or will we enhance or harm human analytic capacity? Will we open a new era of creativity or undermine human creativity? Will AI result in the greatest increase in general productivity and wealth the world has ever seen, or will it be accompanied by the exacerbation of the most profound disparities in wealth and class? Will fake news, counterfeit videos and voices, and bot-driven fake stories compromise our ability to distinguish what is true and undermine national elections? Will we provide soothing robotic friends to children and the elderly, and will that come at the cost of coming to believe that these robotic equivalents are nearly as good as real relationships and so devalue what it means to be human?

These questions are all profoundly important in their own right, and *The Coming Wave* touches on many of them. But the very breadth of the questions suggests something even more significant about AI—that there is a forest we are missing for the particular AI trees. Suleyman explains that to understand AI’s impact, we must grasp the meta-implications of this meta technology—a “general” technology, being disseminated on a scale and at a pace never experienced in human history.

A “General” Technology

As *The Coming Wave* explains, AI is a “general” technology—not only powerful in its own right, but also an accelerant of benefits and risks for almost every other advanced technology. Here is Suleyman:

It’s not just a tool or platform but a transformative meta technology, the technology behind technology and everything else, itself a maker of tools and platforms, not just a system but a generator of systems of any and all kinds. . . . We really are at a turning point in the history of humanity. . . . It won’t be long before AI can transfer what it “knows” from one domain to another, seamlessly, as humans do. AI is far deeper and more powerful than just another technology.

We tend to think of AI as a powerful train coming toward us and focus on its implications for the realities with which we are most familiar—our careers, our finances, our families. Understandable, but to get the big picture it would be better to think of AI as fueling a thousand trains coming toward us and everyone else, in almost every field—some of which are laying their own tracks. Or maybe we tend to think of the AI wave as an army of elephants that will trample some and build new roads for others. It may be better to think of it as millions of ants, many of which are moving and interacting autonomously.

Democratized and Disseminated Power on a Scale and Pace Previously Unseen

The Coming Wave goes on to describe another key element, which distinguishes this “wave” from all prior technological waves: for better and for worse, it is democratizing and disseminating at unprecedented scale and acceleration. In economic terms, the barriers to entry to the use of these powerful, AI-enhanced tools are extremely low, relative to the vast potential externalities they will create. And Suleyman points out that at this disseminated scale, it will be much harder to detect and respond to harms that can arise.

Applying AI to technologically advanced areas from biotech to robotics to nanotech, Suleyman provides compelling descriptions of how the power of AI has already become de-

mocratized—disseminated at relatively low and rapidly decreasing cost into the hands of individuals and small groups, governments, and rogue actors. He readily acknowledges that many, if not most, of these empowered individuals and groups will be responsible, prudent, and well-intended, working for universities, companies, and governments exercising rational oversight, with every incentive to avoid a major mistake.

Even so, as post-COVID19 investigations have demonstrated, these efforts sometimes go horribly wrong: experiments are ill conceived and executed, redundant systems fail, unanticipated results play out, potent experiments leak. And “dual use” threats arise from imprudent researchers or hobbyists without adequate knowledge of what is at stake, for instance, when modifying bacteria or human genetic traits; or from those with nefarious intent and nothing to lose. Suleyman points out that these AI risks are all the more potent because they are fueling and making more accessible other groundbreaking technologies.

If this all seems too sci-fi or conjectural, consider that, in 2016, U.S. Director of National Intelligence William Clapper surprisingly put a new class of genetic engineering techniques like CRISPR/Cas9—described as “word processing for gene editing”—on the annual worldwide threat assessment list of weapons of mass destruction and proliferation. These techniques are “just” tools, not toxins; but they made it to the list because of low barriers to entries and potentially extreme externalities: its broad distribution, low cost, and accelerated pace of development could have “far-reaching economic and national security implications.”

Suleyman describes at some length how AI exacerbates and accelerates these threats. He points out that DNA synthesizers now cost only a few tens of thousands of dollars and are small enough to sit on a bench in your neighbor’s garage, so that people with “grad level training in biology or an enthusiasm for self-directed learning online . . . could soon create novel pathogens far more transmissible and lethal than anything found in nature.” In fact, last year the *Wall Street Journal* reported that *undergrads* at MIT were asked to test whether AI-powered chatbots could be prompted to assist non-experts in causing a pandemic. Within one hour, the chat-

bots produced four potential pandemic-causing pathogens.

The Undertow of the Coming Wave

A corollary concern of the AI risks is how we will respond to them. Suleyman is at pains to describe the risks in a balanced way and always in the context of his inherent optimism about the good AI will bring. But at the same time, he is arguing that it is essential that the world find new means of recognizing and containing the risks, or the harms may well outpace the good, on a scale the world is not prepared to accept. And he says—given that the wave has arrived—we have to do this quickly.

The last section of the book offers an array of containment options, which Suleyman passionately urges. He explains why, in the history of technology, no general technology has ever been put back in its bottle—so the goal must be appropriate containment, which can save us from the Scylla and Charybdis of unacceptable catastrophe or authoritarian government responses. The options he lays out are broad, ranging from regulatory solutions (he does not believe current regulatory structures, as they exist, are able to keep up), to proposing facile and empowered industry-government collaboratives, and a “NASA moonshot” type of initiative to fund innovative efforts to stay ahead of the risks.

You can decide for yourself if you think these are realistic. Suleyman himself seems to doubt that they are, if only because they would require unprecedented cooperation among industry and nations that may not see the need or that are too intent on maintaining an advantage. Failing those solutions, in a way echoing Hawking but also warning against authoritarian solutions, he calls for new world systems with regulatory authority and ability to respond quickly, and collaborations among nations—including among those that do not share a commitment to inalienable rights for their own citizens or for others. Such is his assessment of the risks.

This might be called the coming undertow of the *coming wave*: if things go wrong, and calls arise for exercise of far greater government power to protect us, will people who care about democracy and human flourishing, and the inalienable rights of persons, acquiesce? If we don’t want to, then we must decide what the risks really are, which risks we are prepared to accept,

and which we are not. Prudent responses, where possible, may help us stand against the undertow of sacrificing fundamental and inalienable rights.

One last note: the book touches on, though not at length, another important AI disruption—the impact on our understanding of human flourishing and human dignity. Other writers have expressed concern that we may come to rely on efficient and seemingly “neutral” AI solutions, even though they are not human-defined or necessarily respectful of human life and dignity. As a case in point, Marc Andreessen, AI investor and optimist, who says that AI may save the world, cheerfully argues that “every child will have an AI tutor . . . helping them maximize their potential with the machine version of infinite love.” Sam Altman of Open AI says he does not want AI friends, and does not recommend it, but that may be what we learn to want: “I personally have deep misgivings about this vision of the future where everyone is super close to AI friends, more so than human friends. . . . [A]nd some people are going to build that, and if that’s what the world wants, and what we decide makes sense, we’re going to get that.” That would be a poor, destructive counterfeit for real love. We are already seeing the dramatic results of loneliness and depression accompanying curated, digitized smart-phone relationships, driven by reinforcing algorithms that vie for our attention and engagement. AI-enhanced technologies will take these capabilities to whole new levels, if we are not alert and disciplined enough to resist temptation.

The world needs to hear from those who know about real love, and real personhood. People who believe we are made in the image of God, who live in the hope of the incarnation, redemption, and the knowledge that we are made to be loved and to love, and who understand our responsibility to steward the beautiful world God has entrusted to us, have something of great value to offer a world tempted by these counterfeits. We should be ready to offer that good news because as the coming wave breaks, the world is going to need it even more. A good way to start getting ready might be to read *The Coming Wave*.