

# Henry Kissinger on the Promise and Threats of AI

**November 21, 2023**  
**by Laurence B. Siegel**

At this horrible moment in history, with Israel waging a war on multiple fronts and Ukraine fighting for its independence, it's reassuring that serious thinkers have long been reflecting about national and global security. They are considering the way that AI will interact with military and strategic matters. Thus, *The Age of AI (And Our Human Future)*, by Henry Kissinger, Eric Schmidt, and Daniel Huttenlocher, is keenly relevant even though the book is two years old, an eternity in AI-space.

The book is not a great read. It's wordy, repetitive, and jargon-filled. But, in light of our onslaught of challenges, Chapter 5, "Security and World Order" (pages 135-176), is worthwhile. I begin my review with some comments on that chapter and then move on to the economics of AI and related topics.

## About the authors

Henry Kissinger



[Source](#)

It's a surprise to find that Henry Kissinger has written a book at all - he's 100 years old. It's even more surprising to find that the book is about artificial intelligence, a recent outgrowth of the computer industry, which in turn is younger than Kissinger himself. He was a grown man of 22 when ENIAC, the first general-purpose programmable electronic computer, creaked to life at the University of Pennsylvania.<sup>1</sup>

Eric Schmidt, one of the co-authors of *The Age of AI*, is the former CEO of Google. It was not a surprise to see his name on the cover, along with that of Daniel Huttenlocher, who is the first, and so far, only, dean of the Schwarzman College of Computing at the Massachusetts Institute of Technology.

Much of *The Age of AI* reads like a primer on the topic, addressing questions that all of us have thought about recently. These top-level questions include:

- Is AI really *intelligence*? Or is it a simulacrum of intelligence, a very clever computer program designed, built, and controlled by humans?
- How will we benefit from AI?
- Is AI dangerous?
- What are the consequences of AI for national security?

Eric Schmidt



[Source](#)

The authors also try to answer more profound, but perhaps unanswerable, questions such as, “When AI participates in assessing and shaping human action...what, then, will it mean to be human?” They offer some beautiful words on the topic:

Human identity may continue to rest on the pinnacle of animate intelligence, but human reason will cease to describe the full sweep of the intelligence that works to comprehend reality.... Our emphasis may need to shift from the centrality of human reason to the centrality of human dignity and autonomy.

...but that is a way of saying “we don’t know.”

*The Age of AI* is silent on many of the issues that investors care about: What are the consequences of AI for the economy? Will it set off an economic boom, like those in the past that were fueled by the discovery of a new resource? Will it put many people out of work? How should investors respond? Although *The Age of AI* does not have an economic slant, I’ll address these issues from my own point of view.

Daniel Huttenlocher



[Source](#)

Finally, a question that has popped up recently is: Will AI cause human extinction? No, it won’t. I’ll get back to that, in a suitably mocking tone, after discussing defense and the economy.

## Artificial intelligence, war, and its avoidance

What’s the worst-case scenario? Probably not the silly, but widely circulated, fantasy of an AI that is instructed to make as many paper clips as possible and turns every atom in the universe, including us, into paper clips. No, the ability of AI to start a nuclear war – whether by following an evildoer’s instructions or just by blundering – is AI’s greatest danger. It is the one that Kissinger, Schmidt, and Huttenlocher (henceforth KSH) focus on, so I will too.

### ***Blundering into a nuclear war?***

Here’s a catastrophic scenario that many of us have thought about: An AI-governed missile system is programmed to launch nuclear weapons at an enemy only in the face of overwhelming evidence that the enemy has already begun a nuclear attack. The AI “believes” it has that evidence. But the AI is wrong. There’s no attack. The AI counterattacks anyway. A global thermonuclear war results.

AI is often wrong, as you know from your infuriating experiences with Siri or Alexa. The stakes are low – you try to drive to an address that doesn’t exist, or order chow mein when you meant to order chow fun. But an AI missile system is obviously high-stakes. If you’re thinking “they would build that system with much more

accuracy and many more safeguards,” you’re right – they would.

But our missile detection systems were always built with the intention of being fail-safe. That did not make them fail-safe. They have mistaken a Norwegian spacecraft launch, a harmless Soviet weapons test, and the moon rising over the horizon – *the moon!* – for nuclear attacks. Only human intervention stopped these mistakes from turning into catastrophes. The closest call was on September 26, 1983, when a mid-level Soviet military man, Stanislav Petrov, found that his missile detection system was reporting a U.S. nuclear attack in progress. *Sensing* that there was something wrong with the information, that the U.S. “attack” wasn’t real, Petrov disobeyed orders to retaliate and is widely credited with having saved the world.

AI doesn’t have feelings or hunches or doubts – and it doesn’t have a conscience. It can act too quickly for human intervention and is motivated to do so because any hesitation could result in an AI-directed counterattack that is also too fast for human intervention.

Also, Stanislav Petrov should have won the Nobel Peace Prize.

What do KSH have to say about this?

The range of activities an AI is capable of undertaking...may need to be adjusted so that a human retains the ability to monitor and turn off or redirect a system that has begun to stray...[and] such restraints must be reciprocal... Major powers should pursue their competition within a framework of verifiable limits.

This is weak tea. International agreements, mutual inspections, and willingness to share some secrets make up the peacekeeping framework that Henry Kissinger had a role in building. It worked when only the U.S. and the Soviet Union had nuclear capabilities. *Doveryai no proveryai* – trust but verify.

Those were Cold War safeguards. Today, rogue states and non-state actors could build a nuclear weapon. If you consider North Korea to be a rogue state, one already has. Russia is on the edge of being a rogue state. And AI is much easier to develop than a nuclear weapon. Here’s a grim thought: If an AI in a non-nuclear country, acting under someone’s instructions, *fools* enough people into thinking the country has nuclear capabilities, we could wind up in a “hot” war.

### ***Using AI to do good***

These risks are real, and KSH describe them vividly. But the benign case for AI is even more powerful. We’ve had some sort of AI for generations; we call it “automation,” and it has helped consumers and boosted the efficiency of the economy immensely. Your car’s automatic transmission, which shifts gears at exactly the right moment, relies on hydraulic logic circuits, a form of AI that has been around for 90 years and doesn’t even use electricity. Industrial robots have been working in factories since the 1960s. They used to be fairly stupid, requiring human instruction for every little task; now they “figure out” what needs to be done and how to do it. But they were using AI the whole time. AI is a tool, and any tool can be used for good or ill.

The authors’ best example of using AI to do good is the discovery of a new antibiotic. The drug researchers “developed a ‘training set’,” KSH write, “of two thousand known molecules,” teaching the AI about the chemical properties and ability to inhibit bacterial growth of each. From this training set, the authors note, “the AI...identified attributes that had not...been encoded [and]...that had eluded human conceptualization...” They continue:

When it was done training, the researchers instructed the AI to survey a library of 61,000 molecules...that (1) the AI predicted would be effective as antibiotics, (2) did not look like any existing antibiotics, and (3) the AI predicted would be nontoxic. Of the 61,000, one molecule fit the criteria. The researchers named it halicin – a nod to the AI “HAL” in the film *2001: A Space Odyssey*.

Halicin works against three previously drug-resistant strains of bacteria. While humans could probably have teased out this result eventually, the opportunity cost of the effort involved would have impeded research on other needed drugs. AI performed the work cheaply and well.

### ***Turn it off?***

If AI behaves badly, turn it off! That’s what KSH suggested when discussing military applications, and it makes intuitive sense. What’s wrong with this idea?

Nothing, in principle. It’s exactly what you should do. But turning off a decentralized and intentionally redundant computer system is not easy. You can’t turn the internet off. That is why DARPA, the U.S.

Defense Advanced Research Projects Agency, created the internet in the first place – to make sure that if one part of their global communications system failed, the rest of it would keep humming. So disabling an out-of-control AI that has the nuclear football might be harder than it sounds.

## The economic promise of AI

Enough of this doom and gloom talk. Let's look at the upside.

Almost every economist, financial analyst, and business strategist who has written about AI believes it will increase productivity and thus GDP. Discovery of a new resource or general-purpose technology (AI is both) has always made work easier, workers more productive, and goods and services cheaper in real terms.<sup>2</sup>

How much more productive? An April 5, 2023 Goldman Sachs report said that generative AI could increase global GDP by 7%, or about \$7 trillion. If that's a one-time pop, which its report implies, that is only two or three years' global growth: The global GDP we'd expect in 2100 without AI will arrive in 2097. That would be a disappointment for such a large and potentially risky change in technology.

An increase in the slope or rate of global GDP growth, at least for some years, would be a much better outcome and more likely given the history of other general-purpose technologies such as electricity, air travel, and the internet.<sup>3</sup> The productivity gains from these innovations didn't happen for a while after their discovery, and the commercialization of AI will follow the same trajectory.

McKinsey & Company corroborates Goldman's estimate of a \$7 trillion addition to global GDP but says that the number is for the impact of *generative* AI alone,<sup>4</sup> while *at the same time* forecasting a change in the slope of growth:

Generative AI could enable [additional] labor productivity growth of 0.1 to 0.6 percent annually through 2040... Combining generative AI with all other [automation-related] technologies, work automation could add 0.2 to 3.3 percentage points annually to productivity growth.<sup>5</sup>


To sum up, McKinsey adds the impact of generative and non-generative AI to arrive at an overall "total AI economic potential" of \$17.1 trillion to \$25.6 trillion, one-sixth to one-fourth of global GDP.<sup>6</sup> This is equivalent to adding *another United States* to the world economy.

### ***AI and unemployment***

This increased productivity means that fewer workers are needed to produce the same output. But we do not use the greater efficiency to produce the *same* output! We produce more, or better, or just different goods and services. If you told a 1950s machine operator, who would eventually be replaced by a computer, that his grandchildren would find employment as computer-aided design engineers, social media marketers, or AI researchers, he would have thought you were crazy. But that is what happened. Each technological advance of the last 250 years caused a panic about unemployment, yet the number of people working has grown each time.

If, over the very long term, increased automation took away jobs on net, the unemployment rate would be sky-high and increasing. It isn't. Exhibit 1 shows that while the U.S. population grew from 63 million in 1890 (the first year for which reliable unemployment data exist) to 334 million today, the unemployment rate fell to roughly 4% in every business boom during that long period. Every single time. I call this number the "magic 4%." Obviously, a massive number of jobs were created as we experienced one "job-destroying" technological revolution after another.

## Exhibit 1



Source: Siegel, Laurence B. 2019. *Fewer, Richer, Greener*. Chart based on data sourced and organized by Wikipedia user Peace01234. 1930–2009 data are from Bureau of Labor

Statistics, employment status of the civilian noninstitutional population, 1940 to date. 1890–1920 data are from Romer (1986). 1920–1930 data are from Coen (1973).<sup>7</sup>

### **Labor-force participation**

But hasn't the labor-force participation rate fallen, making the magic 4% more of a statistical artifact than a fundamental economic truth?

You bet it has. Your forebears, if you go back far enough and are not descended from royalty, toiled from sunup to sundown at backbreaking work that barely yielded enough food for those doing it. Children worked about as hard as adults. The labor participation rate was, in round numbers, 100%.

From that inauspicious start, the only possible direction for the labor participation rate to go was down. This represented tremendous progress and is part of what has made modern life enjoyable. There is more to life than work, work, work – and one of the main reasons for developing any technology, including AI, is to make work easier and leisure more available.

But this does not mean that AI lacks an economic downside. It will put some people out of work who cannot realistically be retrained. The amount of labor market churn due to AI could be very large. According to Ethan Ilzetzki and Suryaansh Jain, writing for VoxEU,

The World Economic Forum concluded in October 2020 that while AI would likely take away 85 million jobs globally by 2025, it would also generate 97 million new jobs in fields ranging from big data and machine learning to information security and digital marketing.<sup>8</sup>

This outcome may cause social unrest. It will fail to improve the lives of the hard-core unemployed. And it may stimulate political demands for more social benefits or a basic-income guarantee, further burdening taxpayers including those not yet born.

### **Implications for investors**

Investors are concerned with corporate profits, not economic productivity *per se*. But, while workers and consumers will and should get part of this new bounty, shareholders are able to benefit handsomely. In this sense, the emergence of AI as a new resource or general-purpose technology is strongly bullish. Global equities offer the best medium for taking advantage of this opportunity.

The improvement in efficiency will not arrive all at once, and there will be losses to some companies while others gain. I would not buy "AI stocks" as currently understood. They've already soared. And the stocks that will benefit the most in the long run may be very different from the holdings of AI-oriented ETFs and

funds. Recall the great internet bubble that peaked in 2000. Amazon and Google were “buys” even at those elevated prices if you could manage to hold them through the subsequent bumpy ride, but many of the other bubble stocks are now worthless or nearly so. Meanwhile, quite a few companies formed after the bubble have thrived. At any rate, the benefits of AI may accrue more to companies that use AI effectively – which could be in any industry – than to those that produce the AI technology.

While the upsides of AI may be primarily economic, its potential downsides are not. They are social, cultural, military, psychological, and some would say existential.

## Will AI cause human extinction?

Thankfully *The Age of AI* does not touch on the craze for saying that AI will cause the human species to go extinct. It won't. There are eight billion of us and world population hits a new high every day. A species which does that is not a good candidate for extinction anytime soon.

We are also good at defending ourselves from existential threats, as we can see just by looking at those population numbers. Human extinction caused by AI is about as likely as human extinction caused by climate change – the likelihood rounds to zero. Both are real threats. But killing everyone in the world – that's what “extinction” means – is incredibly hard to do. About 70,000 years ago the world's human population contracted to the point where we can see the residue of the “bottleneck” in the genes of people living today. We don't know what caused the mass dying – an asteroid impact (geologists say there wasn't one), climate change (climatologists assure us that the ice ages were survivable through migration), or maybe a disease. But we're here, so we know that humans survived a terrible catastrophe.

Today, we have an array of tools that our 70,000-year-ago ancestors didn't have: weather satellites, ships, airplanes, instant global communication, antibiotics, and lots and lots of money. We can even nudge a small asteroid out of its path so it doesn't strike Earth. (A large one would be a problem.)

Other than nuclear war, the biggest risk is an AI-engineered bioweapon. But nature engineers those all the time and, despite tragic losses of life, we find ways to fight back and survive. Moreover, we have already developed horrifically destructive bioweapons; luckily, they have not been used. AI-generated ones are unlikely to be worse.

Extinction is the least of my worries.

## Beware of hyperbole!

Almost everyone who writes about AI falls into the trap of overstatement. I thought that the serious-minded KSH team might not, and much of the book is not overstated, but they offer this proposition:

Only rarely...has technology transformed the fundamental political and social structure of our societies... The car replaced the horse without forcing a fundamental shift in our social structure. The rifle replaced the musket, but the general paradigm of conventional military activity remained largely unaltered... But AI promises to *transform all realms of human experience*.<sup>9</sup>

In a tweet, Yann LeCun, an NYU professor and chief AI scientist at Facebook, fired back:<sup>10</sup>

More than stone tools? Fire? Agriculture? Metal technology? Writing? Mathematics? Gunpowder? Printing press? Steam engines? Aviation? Computers?

Not mentioning religion? Feudalism? Theocracy? Rationalism? Capitalism? Communism? Fascism? Liberal democracy?

LeCun has injected a dose of sobriety into this fanciful discussion. Thank you, professor.

### Conclusion

AI is a tool, invented by humans to make various tasks easier. While it will have formidable power, we should not cower before it in the illusion that it is something greater than us that we cannot restrain. The great line in the 1931 movie *Frankenstein*, “You have created a monster, and it will destroy you,” remains a fixture in our minds 92 years later and 205 years after Mary Shelley wrote the novel on which it was based.<sup>11</sup>

Why? Because the idea that we've created something beyond our control is a recurring fear throughout history. The fear has usually – not always – turned out to be unfounded. We can reassure ourselves that, as Mark Mills, a physicist, venture capitalist, and author of *The Cloud Revolution*, said to me, AI is just “really

good automation.” But, like nuclear technology, AI can be used by bad people to do bad things – or by careless people who do bad things without intending to. We had best be prudent.



**Laurence B. Siegel** is the Gary P. Brinson director of research at the CFA Institute Research Foundation, economist and futurist at Vintage Quants LLC, and an independent consultant, writer, and speaker. His books, *Fewer, Richer, Greener and Unknown Knowns*, explore ideas in economics, investing, environmentalism, and human progress. His website is <http://www.larrysiegel.org>. He may be reached at [lbsiegel@uchicago.edu](mailto:lbsiegel@uchicago.edu).



<sup>1</sup> Kissinger immigrated to New York with a group of young people all about the same age, fleeing the Nazis with their parents, in 1938 when he was 15. They all attended the same New York public high school. Most of them became distinguished in one field or another. But I’ve heard secondhand that only Kissinger retained his thick Central European accent – the others spoke unaccented English later in life. The person who told me this thinks that Kissinger’s accent is partly a put-on, designed to make him sound even smarter than the smart youths he immigrated with. (At that time, German was the language of science and academia.)

<sup>2</sup> “Real” conventionally means “adjusted for inflation,” but it can also mean that a quantity is expressed in terms of the number of hours of work needed to produce it. The number of hours can be, but need not be, difficulty-adjusted. See Baily, Brynjolfsson, and Korinek (2023).

<sup>3</sup> In the same publication, Goldman also projects a 1.5% per year boost to productivity over the next decade, a forecast inconsistent with the forecast of a 7% increase in global GDP. A 1.5% annual increase compounds to 16%, not 7%.

<sup>4</sup> They give a range, \$6.1 trillion to \$7.9 trillion.

<sup>5</sup> <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>

<sup>6</sup> The “total potential” is expressed in dollars of added productivity, expressed as an annual rate [FINISH]

<sup>7</sup> Romer, Christina, 1986. “Spurious Volatility in Historical Unemployment Data.” *Journal of Political Economy*, 94(1): 1-37. Coen, Robert M. 1963. “Labor Force and Unemployment in the 1920’s and 1930’s: A Re-Examination Based on Postwar Experience.” *Review of Economics and Statistics*, 55(1): 46-55.

<sup>8</sup> Note that, while these sound like big numbers, each represents just over 1% of the world’s population.

<sup>9</sup> Italics added.

<sup>10</sup> 9:32 a.m., May 13, 2023 on X (the website formerly known as Twitter). LeCun posts as @ylecun.

<sup>11</sup> Spoken by “Dr. Waldman,” played by Edward Van Sloan affecting a heavy Central European accent. Contrary to what some people remember, Frankenstein is not the monster, but the doctor who created him.

***A message from Advisor Perspectives and VettaFi: To learn more about this and other topics, check out our most recent white papers.***