

The National Insecurity of AI

Graham Allison

After a week of intense discussion of the implications of AI for American national security at the Aspen Strategy Group (ASG) Summer Workshop, I am still processing 50 pages of notes. The ASG is unique in its ability to bring together a group of distinguished former officials and private sector leaders for serious, candid, and nonpartisan conversation in a setting that encourages reflection. While high altitude perspectives in strictly off-the-record discussions with know-a-lots clarifies, conversations in which experts are candid about what they don't know are even more valuable. About the impact of AI on American national security, I am still thinking. But for now, I offer six initial takeaways.

1. No One Knows the Risks

First, no one fully knows about the risks posed by applications of AI that could allow significantly larger numbers of evil actors to create bioweapons that could kill hundreds of thousands of people, or even an AGI that could threaten the survival of mankind on earth. Specifically, those advancing the AI frontier recognize the potential of such risks, and those in the policy community are even more inclined to highlight such dangers. But as one cross-examines what the most informed people say, the incandescent truth is no one knows. Specifically: (1) no one knows how likely these risks are and (2) no one knows what to do about them.

2. Consult History

Second, consult history. While recognition that attempting to cope with this unbounded technology in the absence of confident answers about its risks is uncomfortable and even frightening, for perspective it is instructive to “consult history.” The most recent analogous case in which scientists and strategists attempted to manage the advance of a technology that gave governments the capability to kill millions of human beings—and even potentially to extinguish the life of Homo Sapiens on Earth—was the creation of nuclear weapons. For highlights from that history, the article titled “The Path to AI Arms Control” that Henry Kissinger and I coauthored last October in *Foreign Affairs* provides a good starting point.¹

For the first decade after the explosion of atomic bombs destroyed Hiroshima and Nagasaki forcing Japan to surrender in WWII, the smartest scientists and strategists in the world grappled with the question of what this meant for strategy, statecraft, and the future of international order. Read what they said and wrote, and examine what they actually did. It was not only the peace advocates who called on nations to “ban the bomb.” Idealistic dreamers were not the only ones urging the U.S. government to give its nuclear weapons to the UN and empower it to act as a super-national overlord.

One of the toughest and savviest statesmen in American history was Henry Stimson. As the Secretary of War in World War II (a position in which he had also served in World War I), he led the arsenal of democracy to victory over Nazi Germany and Japan. Stimson is rightly revered as one of the “wise men” who won the war and built the peace that followed. But as the war came to an end, Truman asked Stimson for his best idea about what to do with the bomb. And what was Stimson's answer? He recommended that the U.S. share its nuclear monopoly with Stalin's Soviet Union and Churchill's Britain in a great power condominium that would supervise a global nuclear order.

A decade after Hiroshima, the most brilliant scientist of the century, Albert Einstein, and one of the greatest philosophers of the time, Bertrand Russell, led eleven other renowned thought leaders in

issuing the Russell-Einstein Manifesto calling on the governments of the U.S., the Soviet Union, the United Kingdom, and all others to “abolish war.” Riveted by the specter of an existential threat to mankind, they urged governments to “remember your humanity and forget the rest.”² Operationally, that meant accepting whatever outcomes could be achieved through peaceful negotiations— without war or the threat to go to war to defend their nation’s interests.

What that would have meant for the lives of Western Europeans as Soviet troops moved from domination of Eastern and Central Europe to Berlin and beyond gave rise to a heated debate about whether one would rather be “red or dead.” Would Einstein and his colleagues have been prepared to live—or as Jews most likely die—in a Nazi-dominated world rather than risking nuclear war to prevent that from occurring? Certainly the statesmen who developed the U.S. national security posture the U.S. adopted—with nuclear deterrence as its centerpiece—believed that it offered a better alternative. For anyone attempting to find clues in this history for addressing AI today, McGeorge Bundy’s chapter on “The Failure of International Control” in his book, Danger and Survival: Choices About the Bomb in the First Fifty Years is a must read.

So, my second big takeaway is that while those currently at the frontier in thinking about this threat and responses to it are confused, they are no more so than their predecessors seven decades earlier.

3. The U.S. Must Remain the Leader in AI

Third, about the proposition that the U.S. must remain the leader in AI, the American national security community is virtually unanimous. Obviously, China’s Xi Jinping has an analogous ambition for his country. Thus, AI is today and will continue to be at the forefront of the fiercest Thucydidean rivalry history has ever seen. As surely as the U.S. and Soviet Union each did everything in its power to achieve nuclear superiority over the other, the U.S. and China are now competing for supremacy in AI— whatever that means. As Kissinger put it: “Never in history has one great power fearing that a competitor might apply a new technology to threaten its survival and security forgone developing that technology for itself.”³

4. U.S.-China AI Arms Control

Fourth, should the U.S. and China explore possibilities for AI arms control? At this point when American strategists are struggling to get our own heads around the challenge, many argue that talking about these risks to the only other AI superpower—China—in the hope of finding areas of agreement is a fool’s errand. But again, I suggest they consult history. The U.S. and the Soviet Union were among the deadliest adversaries the world has ever seen. In their Manichean struggle, each saw the other’s ambitions as a mortal threat to its survival. Nonetheless, in their early conversations about an unprecedented technology that was rapidly advancing in ways neither understood, they discovered a number of islands of shared interests. For example, while each would have preferred to be a nuclear monopolist, their next best choice was to be duopolists, and after that, oligopolists. Thus, they found common cause in preventing the spread of the most dangerous technology, acting both unilaterally and cooperatively to create what became the nonproliferation regime.

Similarly, while each would have preferred that its adversary not have nuclear weapons, once both nations did, each had an interest in the other government maintaining tight control over the decision to launch its nuclear weapons. Thus, when the U.S. developed “permissive action links” that were essentially electronic locks preventing the use of its nuclear arsenal without the president’s unique authorization, it had an interest in the Soviet Union developing an equivalent for the Soviet nuclear arsenal. After decades of an arms race in which each built arsenals of more than 50,000 warheads, the two negotiated formal arms control treaties limiting the numbers of warheads each deployed. Indeed,

they negotiated agreements binding each not to deploy particular weapons systems that would have been destabilizing as the price for its adversary not deploying its equivalent of that system, for example, in the Antiballistic Missile Treaty.

5. AI Is Not Nuclear Weapons 2.0

Fifth, despite the similarities, AI is not just nuclear weapons 2.0. Indeed, the differences between the challenges posed by AI and those faced by the statesmen who confronted nuclear dangers are more salient than the similarities. Nuclear weapons were invented by governments, owned by governments, and monopolized by governments. Governments maintained the strictest control of the knowledge of how to build nuclear weapons and the technologies required to do so. In contrast, AI is being driven by scientists and technologists, entrepreneurs and companies seeking fame, fortune, and a future that promises huge benefits for mankind. The advanced semiconductors that make possible the “mass compute” required to train LLMs like GPT 4.0 are owned by private companies. They make the key choices about which models to train, which risks to accept as they advance applications of AI, and what information about their models they make public. The gladiatorial competition among companies seeking a gazillion dollar pot at the end of this rainbow is as intense as anything seen in the fiercest competitive days of the Cold War. The speed at which AI is advancing makes the path from the atomic bomb to the exponentially more powerful hydrogen bomb seem anachronistic. And while governments employed the individuals who knew the most about nuclear weapons, those most knowledgeable about AI today are private citizens working for private companies.

Some of those most alarmed about the extreme risks posed by AI have proposed that the U.S. government nationalize the AI companies at the frontier, including OpenAI, Microsoft, Anthropic, Google, Amazon, and others, and require their scientists and engineers to work for the government. In the U.S. today, that is sheer fantasy. The difference between citizens’ confidence in their government in the years after the American government had won the Great War, on the one hand, and Americans’ lack of confidence in their government today, on the other, could hardly be starker. The extent to of what Eisenhower called the “military-industrial complex,” and Senator John McCain amended to label the “military-industrial-congressional complex,” in the current state of American democracy where politicians are constantly raising funds from wealthy individuals and companies whose interests they therefore have to take into account makes such draconian proposals delusional. Absent a major catastrophic failure, for the decisive decade ahead the U.S. government is unlikely to have more success in limiting the advance of AI than it has had on social media companies in the past decade.

6. The Good News

Sixth and finally, some good news. In my formal comments at the Strategy Group Summer Workshop, I proposed that we begin by stepping back and asking about the larger international security order today. That order can be summarized in three numbers: 78, 78, and 9. If one can identify the questions to which each is the answer, he will begin to grasp the big picture of international security over the course of most of our lives.

The first 78 is the number of years since the last great power war. This historically anomalous “long peace” came after two colossal wars, each of which was so destructive that historians found it necessary to create a new category for naming them: World War I and World War II.⁴ This unnatural achievement is truly extraordinary—the result of grace and good fortune, but also wise American statecraft. This unprecedented era of peace is not a finished project—but rather a fragile work in progress that has to be earned every day. It is also an accomplishment that is not likely to be sustained over the next generation. When the next great power war comes, every other priority on our current

lists will be subordinated. When the next great power war comes, it will be the biggest event in recorded history.

The second 78 is the number of years since nuclear weapons were used in war. In 1945 or 1950, had anyone offered to make bets on the likelihood that the next seven decades would not see a single use of nuclear weapons in war, he could have gotten 1000 to 1 odds against that happening. Over the course of almost eight nuclear decades, we have survived a number of close calls—most dangerously the Cuban Missile Crisis. There, in a face-off with the Soviet Union to prevent it placing nuclear-tipped missiles in Cuba, John F. Kennedy estimated the risks of nuclear war as between 1 and 3 and even. Like the first number, the second 78-mile marker does not signal that we have arrived. Just last year, Putin seriously threatened to conduct tactical nuclear strikes on Ukraine. The CIA estimated the odds of a Russian nuclear strike as 50-50.⁵ Fortunately, outstanding statecraft from the Biden administration, and imaginative collaboration with Xi Jinping's China, persuaded Putin that this was a very bad idea. But when the next close call comes, or the one after that, it is not likely that this second stunning achievement will be sustained for the next generation.

Finally, and perhaps most remarkably of all, 9 is the number of nuclear weapons states in the world today. In the 1950s and 60s, national security analysts expected that as nations acquired the technical capability to build nuclear weapons, they would do so. John F. Kennedy predicted that by the 1970s there would be 25-30 nuclear weapons states. Recognizing that in a world of many nuclear-weapon armed states there would be repeated analogs of the Missile Crisis, Kennedy mounted one of the boldest initiatives of American post-war policy in an attempt to prevent that future from happening. What we now celebrate as the “nonproliferation regime,” with its centerpiece, the Non-Proliferation Treaty, in which 185 states have foresworn nuclear weapons, is the result.

Again, this is a fragile achievement. Iran stands on the threshold of nuclear weapons. Two thirds of South Koreans now favor acquiring their own nuclear arsenal to protect themselves against North Korea's growing nuclear arsenal.⁶ Contemplating the possibility of a second Trump administration that might not embrace NATO's Article 5 commitment to defend Europe against Russian aggression, Europeans are beginning to discuss the possibility of a European, or even German, nuclear deterrent.

So, my final big takeaway at this point is that however frightening the risks posed by AI (and they are frightening), and however confused the current conversation is about how to address it (and it is confusing), we should reflect on what earlier generations thought and did. The fact that we have been able to live our entire lives without World War III, without uses of nuclear weapons in war, and without a nuclear anarchy in which nuclear wars would be a recurring feature is a largely unrecognized but almost unbelievable accomplishment. It reflects intelligent, persistent, hard work by successive Democratic and Republican administrations over almost 8 decades. It is an achievement for which we should give thanks every day. And as we face the challenges posed by AI today, it can serve as a source not only of insights but of inspiration.

Graham Allison is the Douglas Dillon Professor of Government at Harvard University and a leading analyst of national security with special interests in nuclear weapons, Russia, China, and decision-making. He was the founding dean of Harvard's John F. Kennedy School of Government, and until 2017, served as director of its Belfer Center for Science and International Affairs, which is ranked the #1 University-Affiliated Think Tank in the world. As Assistant Secretary of Defense under President Clinton and Special Advisor to the Secretary of Defense under President Reagan, he has been a member of the Secretary of Defense's Advisory Board for every Secretary from Weinberger to Mattis. He has the sole distinction of having twice been awarded the Department of Defense's highest civilian award, the Distinguished Public Service Medal, first by Secretary Cap Weinberger and second by Secretary Bill Perry. His first book, *Essence of Decision: Explaining the Cuban Missile Crisis* (1971), ranks among the all-time bestsellers with more than 500,000

copies in print. Dr. Allison's latest book, *Destined for War: Can America and China Escape Thucydides's Trap?* (2017), is a national and international bestseller.

¹ Henry Kissinger and Graham Allison, "The Path to AI Arms Control," *Foreign Affairs*, October 13, 2023, <https://www.foreignaffairs.com/united-states/henry-kissinger-path-artificial-intelligence-arms-control>.

² Ibid.

³ Henry Kissinger and Graham Allison, "The Path to AI Arms Control."

⁴ John Lewis Gaddis, "The Long Peace: Elements of Stability in the Postwar International System," MIT Press, Volume 10, Number 4, 1986, https://muse.jhu.edu/article/446174%252523info_wrap.

⁵ David E. Sanger, "Biden's Armageddon Moment: When Nuclear Detonation Seemed Possible in Ukraine," *The New York Times*, March 9, 2024, <https://www.nytimes.com/2024/03/09/us/politics/biden-nuclear-russia-ukraine.html?smid=nytcore-ios-share&referringSource=articleShare>.

⁶ Lauren Sukin, "The US Has a New Nuclear Proliferation Problem: South Korea," *Bulletin of the Atomic Scientists*, January 19, 2023, <https://thebulletin.org/2023/01/the-us-has-a-new-nuclear-proliferation-problem-south-korea/>.