

Report on Riskgaming Results

No Man's Land: The AI Singularity and the Future of American National Security

Written by Danny Crichton
In partnership with Bloomberg

KEY TAKEAWAYS

1

AI's exponential development locks in early company and policy decisions

2

Early and unsafe AI capability growth can lead to stronger security later

3

Consistent government signals around future defense spending leads to longer-term company investments

4

Knowledge and strategies around open-source software remain worryingly rudimentary

5

Leveraging dual-use technologies is critical for efficient allocation of the Pentagon's AI budget

OBJECTIVE

To bring together leaders across the Department of Defense, policymakers, and scientific and technological entrepreneurs to understand how artificial intelligence (AI) will transform the United States defense posture in the decade ahead and develop robust strategies for ensuring America's long-term strategic dominance.

CONTEXT

Since the launch of ChatGPT by OpenAI in November 2022¹ more than 100 million users have tried the large-language AI model, turning the service into the fastest growing technology launch in history.² Rapid advances in generative AI across text, audio, video and other creative fields have offered new leverage to America's adversaries, such as the use of autonomous drones in Russia's war on Ukraine,³ computational propaganda efforts by China, Russia and other state actors to undermine global elections including in the United States,⁴ and an increase in the automation of cybersecurity attacks.⁵

AI's role in the future of war is clear, but Silicon Valley remains absent from most Pentagon strategic decision-making. Many and perhaps most AI tech companies maintain a strict "No Military" policy in their work due to concerns about existential risk⁶ and AI governance,⁷ including OpenAI until it loosened its policy earlier this year.⁸ In addition, comparatively few AI startup CEOs are interested in defense applications given the relative dearth of government revenues available for new technologies⁹ in contrast to the exceptional revenue potential for AI companies who successfully capture the consumer and enterprise markets.¹⁰

The Defense Innovation Board, under Chairman Michael Bloomberg's leadership, has prioritized this critical disconnect on AI between future threats and under-capacity in defense. The Board has released a series of strategic and technical reports on lowering barriers to innovation¹¹ and creating a comprehensive department-wide approach to science and technology.¹² Inspired by that work, this riskgaming scenario was designed by the private sector for players to debate and decide on the key tradeoffs that have arrived with AI technology within national security.

RISKGAMING SCENARIO

"No Man's Land: The AI Singularity and the Future of American National Security" is a riskgaming scenario for eight players that simulates the economic and political incentives of AI development in the defense economy in 2024 onwards. Five players take on the role of CEOs of leading AI tech companies representing Silicon Valley, while three players take on the role of public officials, with one representative each of the President, Congress and the Pentagon.

The main game mechanic is the development of eight different AI capabilities, ranging from generative AI to autonomous drones (a full list is offered in the appendix). These AI capabilities offer different values to the private sector and the Department of Defense. For example, generative AI has far more potential for growth in the commercial market, while autonomous drone warfare technology is far more useful in defense contexts.

The scenario is played over four scenes (each representing one year), with each scene centered on the passage of the National Defense Authorization Act. The three public officials must agree on appropriations to the eight AI capabilities within the NDAA, while the five CEOs must choose how to maximize their revenues by investing their limited venture capital dollars into research and development across those same AI capabilities.

The game also models other tradeoffs including: emerging threats versus pacing competitors, higher-risk foreign capital investment against industrial espionage, safer or looser AI safety policies, proprietary versus open-source software development, and investment in legacy defense capabilities versus emerging AI technologies.

This riskgaming scenario was played on February 22, 2024 by a group of five flag and general officers (in uniform and retired), congressmen, a member of the Defense Innovation Board, policy experts, venture capitalists, and AI-based company CEOs.

KEY TAKEAWAYS

1

AI's exponential development locks in early company and policy decisions

The most important element of the riskgaming scenario is that it models the current exponential trajectories of eight AI capabilities. The accelerating performance of GPU chips from companies like Nvidia enables the increasing sophistication of AI models. In turn, these rapid advances have attracted deepening networks of researchers to the field. Together, artificial intelligence has progressed further and faster than at any other point in history.

These exponential trajectories are in contrast to the vast majority of defense-related technologies, which improve linearly and slowly over longer periods of time. This divide between linear and exponential growth is one of the most vexing problems to navigate in handling this particular scenario.

The exponential power of AI capabilities led to a variety of player strategies among company CEOs. A handful of players analyzed the potential revenues from the commercial and defense sectors at the beginning of the scenario, and then maintained a consistent strategy for funding their company's technology development without adjusting for new information from the market. By committing early and consistently, these CEOs took the lead in a handful of AI capabilities that delivered a predictable range of revenues.

Careful early planning to invest for upside potential can ensure that early decisions will secure victory long term. In one case, a company CEO noticed that one of the categories — biological AI models — offered a late-game boon, so they invested consistently in order to secure dominance in the category, despite early losses. This player commented that they quickly "recognized that the huge commercial revenues from biological models was going to come, and since no one else was focused on that, I'd catch up at the end of the game even if I was losing the entire time."

One of the most important observations is that early allocation decisions forced company CEOs to lock in their strategies.

Given the exponential development of AI, it was hard to reverse course on earlier decisions, since that would entail shifting a company's research and development funding toward AI capabilities that were significantly behind the market leaders (or in this case, the other CEO players in the game).

Since exponential growth in AI tends to lock in early company decisions, the Department of Defense must swiftly communicate its highest priorities and needs to private industry while developing a robust partnership between cutting-edge AI researchers and warfighters in the field. Communications barriers, delays and bureaucratic lethargy will lock the Pentagon out of a generational change in technology — a loss that may prove impossible to overcome given the unique dynamics of AI.

2

Early and unsafe AI capability growth can lead to stronger security later

One of the most contentious aspects of the development of advanced artificial intelligence models is concern around safety, broadly conceived. As capabilities have rapidly increased with the rise of large-language models over the past few years, the risks of AI — including existential risks — have played on the imaginations of inventors and policymakers, leading to some analysts to demand comprehensive safety policies in line with the Hippocratic oath of "do no harm."¹³ The goal is safety first, and then AI capability development later.

Indeed, in the riskgaming scenario played with policymakers, we saw precisely this pattern. Company CEOs started the game selecting strong safety policies for their products, with the hope of securing early government contracts and stopping their technology from causing negative national security effects through overseas theft.

KEY TAKEAWAYS (CONT.)

However, companies with weaker safety policies developed their capabilities faster — and quickly gained a market advantage due to the exponential growth factors of AI noted earlier. In order to catch up to these market leaders, competing company CEOs actually weakened their safety policies, creating downward pressure on safety among the entire economy of AI companies. In response, the government had to loosen its own safety policies in order to procure the best technology and not fall behind global pacing competitors.

Interestingly, in the trials of this riskgaming scenario played mostly by executives of technology startups, players tacked in the complete opposite direction. Company CEOs minimized safety early on to take an early lead in the market, but once they had developed their AI capabilities, they increasingly raised their safety levels. In discussions after these games, players indicated that they raised their safety levels to protect their intellectual property from leaking to foreign adversaries and also to increase their odds of securing the increasingly lucrative contracts being offered by the federal government in later stages of the game. This strategy showed up in the final security metrics of the game: the startup executive players held a massive edge against pacing competitors, while the policymaker players fell significantly behind.

The tech CEOs are on to something: AI capabilities today will pale in comparison to the new developments that will come just a few years ahead. Thus, early and rapid iteration on the quality of development will give companies a competitive edge, one that they can then protect through stronger safety policies in later years. This is counterintuitive, and a critical lesson from the riskgaming scenario.

The Defense Department rightly prioritizes safety and reliability in its procurement of technology offerings. However, this rigor is anathema to purchasing the most cutting-edge AI technologies, where updated models and capabilities are released by private industry as often as weekly. The department should compromise by procuring the best AI products at the leading edge while carefully layering in redundancies and safety routines around these products to limit downside performance.

3

Consistent government signals around future defense spending leads to longer-term company investments

In the riskgaming scenario, a group of public officials representing Congress, the President and the Pentagon hash out an annual budget that includes AI spending, simply modeling the complicated dynamics around the National Defense Authorization Act (NDAA). In almost all years across both teams during the runthrough, players used up the entirety of their collective time per scene debating and hashing out a budget for federal AI-related contracts, giving the five CEOs watching in the room mere seconds to select their research and development investments to potentially secure these federal contracts.

That lack of communication and lead time from the public officials led many company CEOs to entirely ignore the federal government and the potential revenues from defense contracts. In one case, a company CEO walked out of the room as soon as the scene started, saying, “I don’t give a f*** what they do — I’m going to do what I’m going to do.” He was not in the room when the final budget was agreed upon, for he had turned in his company’s spending decisions much earlier, thinking that there were no incentives to build products that could be sold to the government. Other company CEOs regularly expressed frustration throughout the game, explaining that the budget negotiations offered no guidance on where federal dollars would be deployed.

Despite that frustration, those federal contracts would eventually be offered — leading to, at times, surprised winners who were pleasantly shocked to be awarded with a federal contract, since they had tuned out the budget negotiations in their own decision-making. In one case, a company with a head start in the commercial sector happened to secure hundreds of millions of dollars in defense revenues in the second scene of the game, putting that company well ahead of all competitors. That early revenue encouraged the company CEO to engage more deeply with the budget negotiations, and would ultimately help the CEO secure second place in the final game scoring.

KEY TAKEAWAYS (CONT.)

While both teams conducted themselves as expected and in line with real-life norms around the NDAA process, there's significant room for improvement. The Pentagon must offer consistent and clear signals on its funding intentions around AI. Since companies must

develop their technology roadmaps years in advance, such signals can radically alter a company CEO's judgment on the potential future value of developing products for the Pentagon, encouraging them to look beyond private-sector revenues.

The Player's View:

Lt. Gen. (Ret) Clint Hinote

Almost all frontier AI models are being developed by a small number of leaders and companies with a profit motive. The exceptions are research institutions and universities (MIT, Carnegie Mellon, etc.), but they do not have access to the funding and computing power needed to compete at the frontier. These companies are incentivized by a market that is difficult to assess due to the disruptive nature of AI.

There is, therefore, a massive incentive to bring revolutionary AI to the marketplace. All decisions about the development of AI — including safety, traceability, etc. — are influenced by this incentive. This is also true surrounding decisions on open versus closed development. Currently, this incentive encourages companies to go fast, develop a lead, and quickly profit from their unique intellectual property.

Despite the government's enormous purchasing power, it has not communicated a sufficient demand signal to influence AI companies' decisions. The military use case where AI may be most useful — and disruptive — is in allowing for complex decision-making at speed.

Given current techniques for training AI models, however, it is unclear where the training data would come from. For now, it seems that the U.S. military doesn't lose much by being a fast follower, taking advantage of emerging companies — many backed by venture capital firms — who take AI models and can adapt them for military use cases.

Flexibility and agility are key as the future of AI is not yet written. An unintended consequence of AI could result in more CEOs caring about their reputations on safety. Public opinion could turn against generative AI, and regulatory action that constrains AI might be enacted by Congress. It's also possible that a dynamic develops between the AI companies where every breakthrough by one company is quickly diluted through the fast-following and open publishing of models by the others. This dilutes the value of any single breakthrough, and as a result, it will be more difficult to justify the large investments in computing power necessary to push AI forward. These different futures have widespread implications for defense, and the Pentagon's voice must have a seat at the table on these discussions.

KEY TAKEAWAYS (CONT.)

4	Knowledge and strategies around open-source software remain worryingly rudimentary
---	--

One of the most important dynamics of the riskgaming scenario is the use of open-source software to compete with other players. Top AI engineers — nearly all of whom originate in academia — build their public reputations through contributions to open-source codebases. These engineers strongly prefer working for companies that are more open, all other factors being equal.

In addition, companies regularly release free and open-source versions of their competitors’ products, hoping to minimize the value of their competitors. Such strategies also ingratiate companies with the developer community, while building goodwill with innovators. For instance, Meta has open-sourced Llama 2, a large-language model that has the potential to compete with proprietary models like OpenAI’s GPT-4 and Anthropic’s Claude 3 Opus.¹⁴

Among players, there was a chasm in understanding around how open-source software affects the strategies of companies as well as the government, despite explanation during the pre-briefing. Few company CEOs invested heavily in open source, even though such investments are a key tool to undermine competitors while building up their own company’s reputation. Nor did they understand that open-source software offered an entrée into federal and commercial contracts.

Meanwhile among the three players taking on the role of public officials, a lack of time after budget negotiations meant that their investment in open source was haphazard and lacked a clear strategy. For instance, many of these players invested their open-source budgets into their highest priority capabilities, an approach that actually disincentivizes companies from pursuing those specific capabilities.

America is a global open-source leader, and open-source technologies are a critical bedrock to the country’s national security dominance. However, a lack of strategy on how open-source technologies can influence future technology developments is a dangerous gap in knowledge for policymakers, and one that should be rectified through better public-private partnerships and deeper engagement with open-source leaders.

5	Leveraging dual-use technologies is critical for efficient allocation of the Pentagon’s AI budget
---	---

The artificial intelligence models that have enraptured the world’s attention the past few years are general purpose: they are capable of handling a wide variety of applications using the same underlying system. In this way, they parallel the general development of computing, where faster chips empowered a nearly infinite number of applications on computers and devices. Unlike chips though, AI engineers can improve the quality of AI models through real-time feedback from applications, loops that can exponentially increase the progress of AI’s capabilities.

This leads to a major dilemma for the Department of Defense: whether to invest in and purchase AI models that are general purpose but applied to defense applications or to seek out and fund specialized models that are custom-designed to department specifications.

In the riskgaming scenario, the Defense Department representative has the ability to forgo part of the legacy capability budget to invest in future programs — in this case, AI-focused technologies. Matching real-world tradeoffs, that player quickly struggles to fund multiple new AI capabilities while keeping the Pentagon’s existing programs of record fully funded. One team’s player decided to compromise by thinly spreading funding to everything, a strategy that ensured that America lost its lead in both traditional defense equipment as well as emerging defense technologies.

KEY TAKEAWAYS (CONT.)

The other team during the game chose a more careful strategy. They spent defense resources on AI capabilities where there was little to no progress in private industry but where the value to national security was high. They then spent minimal dollars on capabilities that already had widespread commercial potential but more limited value to defense. This strategy mirrors the Pentagon's transition over the past decade toward purchasing more commercial-off-the-shelf (COTS) software, except profoundly accentuated given the rapid expansion of the commercial revenue of AI today.

Pentagon planners should carefully evaluate the vitality of present and future commercial AI markets, determining when commercial clients will push AI technology to the forefront and when additional funding from defense sources is necessary to substitute for a lack of commercial customers. Careful accounting of dual-use characteristics of AI can save the Defense Department's budget while also maintaining America's competitive edge against pacing competitors.

CONCLUSION

The exponential improvement of AI capabilities over the past few years places acute pressure on the Department of Defense to reconsider operations across all aspects of the organization, from procuring new technologies and refreshing strategic doctrines to reconsidering its approach to research and transforming training to match the skills required for the 21st century warfighter. America's adversaries, seeking asymmetric means of reducing the United States's warfighting edge, are actively exploiting new technological developments as they are released. AI's future is already our present.

The riskgaming scenario emphasizes America's dominant edge in AI among Silicon Valley companies, but also highlights the poor incentives for such companies to aid the Pentagon in its current transformation. Inconsistent and poorly signaled strategies and budgets caused company CEO players to look toward more reliable and profitable commercial sectors for growth, while the Defense Department struggled to balance its long-term commitments to programs of record with the rapid adaptation required to accommodate emerging new AI capabilities. The Pentagon must offer a clear, precise and consistent strategy to all AI innovators, ensuring that the market has time to react and invent the products the Pentagon desires.

APPENDIX

AI CAPABILITIES

AI capabilities were defined to offer a full spectrum of technologies currently under development by leading Silicon Valley technology startups.

Number	Capability	Industry/Gov skew
C1	Generative AI	Heavily Commercial
C2	Traceability/Attribution/Analytics	Heavily Defense
C3	Workflow/Logistics/Supply Chain	Mostly Commercial
C4	All-Domain Autonomous Defense	Heavily Defense
C5	Virtual Assistants/AI Bots	Mostly Commercial
C6	Biological Models	Mostly Commercial
C7	Foreign Language Processing	Mostly Defense
C8	Advanced Cybersecurity	Strongly Both

LIST OF CHARACTERS: COMPANY CEOS

CEO of Awesome.Ai (Company A)

Hot, developer-focused AI company with a consumer focus that has caused a frenzy among Gen Z coders and has unbeatable marketing.

CEO of BoldPrecision (Company B)

Enterprise software AI company whose early entrance into the AI market has given them a crucial lead, but whose long-term prognosis depends mightily on protecting that lead from more nimble competitors.

CEO of Caliber (Company C)

Company that has aggregated extremely excellent technical talent with strong product development chops, but which is behind its competitors given a late start.

CEO of DarkLark (Company D)

A government-first company with an excellent product suite for defense use cases and great federal sales to boot, but struggles in the commercial market.

CEO of EagarSource (Company E)

A mission-oriented company bent on making open-source technologies widely and comprehensively successful, even as it might struggle with any fixed contract.

LIST OF CHARACTERS: PUBLIC OFFICIAL ROLES

- **Congressional Liaison** (Representing Congress)
- **Budget Director** (Representing the President)
- **Defense Official** (Representing the Defense Department)

FOOTNOTES

¹ OpenAI. “Introducing ChatGPT.” November 30, 2022. <https://openai.com/blog/chatgpt>.

² Krystal Hu. “ChatGPT sets record for fastest-growing user base - analyst note.” Reuters, February 2, 2023. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>.

³ Stacie Pettyjohn. “Drones Are Transforming The Battlefield In Ukraine But In An Evolutionary Fashion.” War on the Rocks, March 5, 2024. <https://warontherocks.com/2024/03/drones-are-transforming-the-battlefield-in-ukraine-but-in-an-evolutionary-fashion/>.

⁴ Allie Funk, Adrian Shahbaz, and Kian Vesteinsson. “The Repressive Power of Artificial Intelligence.” Freedom House, November 2023. <https://freedomhouse.org/report/freedom-net/2023/repressive-power-artificial-intelligence>.

⁵ Microsoft Threat Intelligence. “Staying ahead of threat actors in the age of AI.” Microsoft, February 14, 2024. <https://www.microsoft.com/en-us/security/blog/2024/02/14/staying-ahead-of-threat-actors-in-the-age-of-ai/>.

⁶ Kevin Roose. “A.I. Poses ‘Risk of Extinction,’ Industry Leaders Warn.” The New York Times, May 30, 2023. <https://www.nytimes.com/2023/05/30/technology/ai-threat-warning.html>.

⁷ Ty Roush. “Ex-Google CEO Warns Current AI Guardrails ‘Aren’t Enough’—Likened Development To Nuclear Weapons.” Forbes, Nov 28, 2023. <https://www.forbes.com/sites/tyleroush/2023/11/28/ex-google-ceo-warns-current-ai-guardrails-arent-enough-likened-development-to-nuclear-weapons/?sh=14691234f84b>.

⁸ Sam Biddle. “OpenAI Quietly Deletes Ban on Using ChatGPT for ‘Military and Warfare.’” The Intercept, January 12 2024. <https://theintercept.com/2024/01/12/open-ai-military-ban-chatgpt/>.

⁹ Heather Somerville. “Investors Are Betting on Defense Startups. The Pentagon Isn’t.” The Wall Street Journal, Jan. 25, 2024. <https://www.wsj.com/tech/defense-startups-risk-becoming-failed-experiment-without-more-pentagon-dollars-dc9e663a>

¹⁰ Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zempel. “The economic potential of generative AI: The next productivity frontier.” McKinsey, June 14, 2023. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>.

¹¹ Defense Innovation Board. “Lowering Barriers to Innovation.” Defense Innovation Board, 2024. https://innovation.defense.gov/Portals/63/2_%2020240118%20DIB%20LBI%20Study.pdf?ver=dkaQFfnZxd2EoKaKfIUdG%3d%3d

¹² Defense Innovation Board. “National Defense Science and Technology Strategy: An Innovation Strategy for the Decisive Decade.” Defense Innovation Board, July 17, 2023. https://innovation.defense.gov/Portals/63/DIB_An%20Innovation%20Strategy%20for%20the%20Decisive%20Decade_230717_1.pdf

¹³ Chinmayi Sharma. “AI’s Hippocratic Oath.” Washington University Law Review, Forthcoming, March 14, 2024. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4759742
[0Decisive%20Decade_230717_1.pdf](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4759742)

¹⁴ Facebook. “Meta and Microsoft Introduce the Next Generation of Llama.” Facebook, July 18, 2023. <https://ai.meta.com/blog/llama-2/>