HARPER'S MAGAZINE

March 2024 Issue

[Letter from Washington]

The Pentagon's Silicon Valley Problem

by Andrew Cockburn

How Big Tech is losing the wars of the future

Illustrations by Yoshi Sodeoka

Three months before Hamas attacked Israel, Ronen Bar, the director of Shin Bet, Israel's internal security service, announced that his agency had developed its own generative artificial intelligence platform—similar to ChatGPT—and that the technology had been incorporated quite naturally into the agency's "interdiction machine," assisting in decision-making "like a partner at the table, a co-pilot." As the Israeli news site *Tech12* explained in a preview of his speech:

> The system knows everything about [the terrorist]: where he went, who his friends are, who his family is, what keeps him busy, what he said and what he published. Using artificial intelligence, the system analyzes behavior, predicts risks, raises alerts.

Nevertheless, Hamas's devastating attack on October 7 caught Shin Bet and the rest of Israel's multibillion-dollar defense system entirely by surprise. The intelligence disaster was even more striking considering Hamas carried out much of its preparations in plain sight, including practice assaults on mock-ups of the border fence and Israeli settlements—activities that were openly reported. Hamas-led militant groups even posted videos of their training online. Israelis living close to the border observed and publicized these exercises with mounting alarm, but were ignored in favor of intelligence bureaucracies' analyses and, by extension, the software that had informed them. Israeli conscripts, mostly young women, monitoring developments through the ubiquitous surveillance cameras along the Gaza border, composed and presented a detailed report on Hamas's preparations to breach the fence and take hostages, only to have their findings dismissed as "an imaginary scenario." The Israeli intelligence apparatus had for more than a year been in possession of a Hamas document that detailed the group's plan for an attack.

Well aware of Israel's intelligence methods, Hamas members fed their enemy the data that they wanted to hear, using informants they knew would report to the Israelis. They signaled that the ruling group inside Gaza was concentrating on improving the local economy by gaining access to the Israeli job market, and that Hamas had been deterred from action by Israel's overwhelming military might. Such reports confirmed that Israel's intelligence system had rigid assumptions of Hamas behavior, overlaid with a racial arrogance that considered Palestinians incapable of such a large-scale operation. AI, it turned out, knew everything about the terrorist except what he was thinking.

Such misplaced confidence was evidently not confined to Israeli intelligence. The November/December issue of *Foreign Affairs* not only carried a risibly ill-timed boast by national security adviser Jake Sullivan that "we have de-escalated crises in Gaza," but also a paean to AI by Michèle Flournoy. Flournoy is a seasoned denizen of the military-industrial complex. The undersecretary of defense for policy under Barack Obama, she transitioned to, among other engagements, a lucrative founding leadership position with the defense consultancy WestExec Advisors. "Building bridges between Silicon Valley and the U.S. government is really, really important," she told *The American Prospect* in 2020. Headlined AI IS ALREADY AT WAR, Flournoy's *Foreign Affairs* article invoked the intelligence analysts who made "better judgments" thanks to AI's help in analyzing information. "In the future, Americans can expect AI to change how the United States and its adversaries fight on the battlefield," she wrote. "In short, AI has sparked a security revolution—one that is just starting to unfold." This wondrous new technology, she asserted, would enable America not only to detect enemy threats, but also to maintain complex weapons systems and help estimate the cost of strategic decisions. Only a tortuous and hidebound Pentagon bureaucracy was holding it back.

Lamenting obstructive Pentagon bureaucrats is a trope of tech pitches, one that plays well in the media. TECH START-UPS TRY TO SELL A CAUTIOUS PENTAGON ON A.I. ran a headline in the *New York Times* last November over a glowing report on Shield AI, a money-losing drone company for which Flournoy has been an adviser. Along with a review of the company's drone, the story cites:

> the many hurdles that the new generation of military contractors face as they compete for Pentagon funding against the far bigger and more entrenched weapons makers that have been supplying the military for decades.

But the notion that the Pentagon is resistant to new technologies is hardly fair—it reportedly funds at least 686 AI projects, including up to $9 billion for a Joint Warfighting Cloud Capability contract awarded in 2022 to a slew of major tech companies destined to, per one Pentagon official, "turbocharge" AI solutions. Another AI project, Gamechanger, is designed to enable Pentagon employees to discover what their giant department actually does, including where its money goes. A press release from the Joint Artificial Intelligence Center from early 2022 that celebrates Gamechanger's

inauguration noted "28 Authoriative [*sic*] Sources" and quoted a senior Pentagon accountant's excitement about "applying Gamechanger to gain better visibility and understanding across our various budget exhibits." Even so, the Pentagon failed to pass a financial audit in 2023, for the sixth year in a row.

Artificial intelligence may indeed affect the way our military operates. But the notion that bright-eyed visionaries from the tech industry are revolutionizing our military machine promotes a myth that this relationship is not only new, but will fundamentally improve our defense system—one notorious for its insatiable appetite for money, poorly performing weapons, and lost wars. In reality, the change flows in the other direction, as new recruits enter the warm embrace of the imperishable military-industrial complex, eager to learn its ways.

The belief that software can solve problems of human conflict has a long history in U.S. war-making. Beginning in the late Sixties, the Air Force deployed a vast array of sensors across the jungles of Southeast Asia, masking the Ho Chi Minh trail along which North Vietnam supplied its forces in the south. Devised by scientists advising the Pentagon, the operation, code-named Igloo White, and designed to detect human activity by the sounds of marching feet, the smell of ammonia from urine, or the electronic sparks of engine ignitions, relayed information to giant IBM computers housed in a secret base in Thailand. The machines were the most powerful then in existence; they processed the signals to pinpoint enemy supply columns otherwise invisible under the jungle canopy. The scheme, in operation from 1967 to 1972 at a cost of at least hundreds of millions a year, was a total failure. The Vietnamese swiftly devised means to counter it; just as Hamas would short-circuit Shin Bet algorithms by feeding the system false information, the Vietnamese also faked data, with buckets of urine hung in trees off the trail, or herds of livestock steered down unused byways, which were then dutifully processed by the humming computers as enemy movements. Meanwhile, North Vietnamese forces in the south were well supplied. In 1972, they launched a powerful offensive using hundreds of tanks that went entirely undetected by Igloo White. The operation was abandoned shortly thereafter.

The IBM System 360 computers at the center of Igloo White were a prominent icon of the industry we now call Silicon Valley. Born of the electronics industry that helped secure victory in World War II, this sector flourished under Pentagon patronage during the Cold War. The development of integrated circuits, key to modern computers and first produced by Texas Instruments in 1958, was powered by an avalanche of defense dollars, and initially deployed in the guidance system for the Minuteman II intercontinental nuclear missile. Beginning with personal calculators, the microchip revolution eventually found a commercial market severed from the umbilical cord of government contracts, generating an industrial culture, eventually both physically and spiritually centered south of San Francisco—far removed from the Pentagon parent that had spawned it.

As America's manufacturing economy gradually declined from the Eighties on, its Rust Belt heartland increasingly studded with decaying industrial cities, the digital economy grew at an exponential rate. Less than a month after its stock market debut in December 1980, Apple Computer was worth more than Ford. It was an industry happy to be independent of the irksome constraints of government contracting. Apple's Macintosh personal computer, the young company imagined, would free citizens from the Orwellian world of government control "by giving individuals the kind of computer power once reserved for corporations." But even though the original freewheeling iconoclasts of the tech industry saw themselves as "cowboys," "rebels," and "revolutionaries," in the words of historian Margaret O'Mara, the divorce from defense was never absolute. The internet, hailed as a liberating technology, grew out of ARPANET, which had been developed by the Pentagon's Advanced Research Projects Agency. According to Yasha Levine in *Surveillance Valley,* the proto-internet was deployed almost immediately to collect information on the antiwar movement. Google founders Larry Page and Sergey Brin's early work at Stanford was funded in part by the same Pentagon agency, which by now had added "Defense" to its name to become DARPA. And Google Earth started as Keyhole EarthViewer, a mapping system partly funded by the CIA's venture capital offshoot In-Q-Tel, which was eventually acquired by Google.

Meanwhile, the Pentagon never lost sight of the unfulfilled dream of Igloo White: that computing power could make it possible to control the battlefield. Following the American retreat from Vietnam, vast sums were poured into Assault Breaker, in which powerful airborne radars would peer deep behind Soviet lines in Eastern Europe. It, too, abjectly failed its tests—the system could not distinguish tanks from cars, or from trees blowing in the wind—and the project was canceled the following decade. Yet military optimism was undimmed; senior commanders broadcast the notion of "netcentric warfare," and their aspirations found fruit in such projects as the Future Combat Systems program, which linked sensors and weapons via high-powered processors to strike targets so effortlessly that, or so its proponents claimed, it would no longer be necessary to install defensive armor on tanks. But after consuming almost $20 billion of taxpayer dollars, it came to nothing. Same with the Department of Homeland Security's Secure Border Initiative

Network, marketed as a "virtual fence" equipped with computer-linked radar, cameras, and other surveillance sensors to detect intruders, which was canceled in 2011. Boeing had been a prime contractor on both of these baroque endeavors, a sclerotic contrast to the fast-paced dynamism of Silicon Valley. Surely the flourishing entrepreneurial culture of the West Coast could succeed where the old guard had failed.

Peter Thiel certainly thought so. A former securities lawyer, he had garnered his initial fortune as a co-founder of PayPal and grew vastly richer due to an early investment in Facebook. A professedly conservative libertarian contemptuous of "hippies" who had "[taken] over the country," he was determined, in the words of his biographer Max Chafkin, "to bring the military-industrial complex back to Silicon Valley, with his own companies at its very center." Founded in 2003 and based on technology developed for fraud detection at PayPal, Thiel's software company Palantir used data to detect and communicate patterns in simple visual displays. Despite proclaiming that "it's essential to preserve fundamental principles of privacy and civil liberties while using data," the CIA, via In-Q-Tel, was an early investor, and many of Palantir's initial contracts reportedly came from the intelligence community. As proof of its worth to such customers, it was claimed that Palantir had a major role in detecting not only the Chinese espionage network GhostNet in 2009, but the secret lair of Osama bin Laden in 2011. There is no evidence that these boasts were justified, but the claims entranced the media and attracted corporate customers. In the face of high-level military reluctance, Palantir marketed the wonders of its technology to mid-level Army officers and engaged cooperative lawmakers in Congress to lobby on its behalf.

Proponents of the Silicon Valley approach found friendly reception elsewhere in the upper tiers of the Pentagon—especially among acolytes of Andrew Marshall, the venerated former director of the Office of Net Assessment, an internal Pentagon think tank. One of them, embedded in Washington's defense archipelago, was a former Marine named Robert O. Work. Appointed under Obama as undersecretary of the Navy in 2009 and rapidly promoted, Work was an ardent advocate of the notion that the United States was losing its technological lead, previously assured by its superiority in nuclear weapons, and then by its precision-guided weapons. Now the Chinese and Russians, he warned darkly, had caught up, endangering America's military dominance. This threat, he believed, was rendered more urgent by caps on defense spending promised in the Budget Control Act of 2011. The danger could only be warded off by adopting, among other things, aerial and naval unmanned systems and AI-enabled battle networks. These were to be found in Silicon Valley.

In 2014, Secretary of Defense Chuck Hagel unveiled the Defense Innovation Initiative, overseen by Work, which, Hagel said, would "actively seek proposals . . . from those firms and academic institutions outside DoD's traditional orbit." Hagel's successor, Ash Carter, made repeated pilgrimages to tech enclaves, shedding his tie and delivering boosterish homilies on the rewards of the collaboration. "They're inventing new technology, creating prosperity, connectivity, and freedom," he told a reporter following an early visit. "They feel they too are public servants, and they'd like to have somebody in Washington they can connect to." In 2015, Carter launched the Defense Innovation Unit Experimental ("Experimental" would later be dropped), with an office in Mountain View, near the headquarters of Google and Apple. The Department of Homeland Security had the same idea, setting up a satellite office nearby to "cultivate relationships with Silicon Valley technology companies, particularly non-traditional performers to help them understand DHS's challenges." In 2016, Carter gave the budding relationship potent bureaucratic heft, convening a Defense Innovation Advisory Board populated with heavyweights from the tech and military worlds tasked with making recommendations directly to the secretary of defense. Heading the board was Eric Schmidt, former executive chairman and CEO of Google. Schmidt's role soon grew larger, to head of the National Security Commission on Artificial Intelligence, with Work as his deputy. They were charged with advancing "the development of artificial intelligence, machine learning, and associated technologies to comprehensively address the national security and defense needs of the United States."

Among the aims of Work's AI initiative was a means to distill the vast amounts of information sucked up by satellites, phone intercepts, emails, and drones sitting in intelligence and military data banks into something accessible. In 2017, Google secured a contract for Project Maven, instituted by Work, aimed at accelerating "DoD's integration of big data and machine learning." Project Maven's initial goal was to develop processing tools for the unceasing torrent of drone footage in order to identify targets. The contract, according to internal emails leaked to *Gizmodo*, specified that Google's involvement was not to be revealed without the company's permission. "Avoid at ALL COSTS any mention or implication of AI," Fei-Fei Li, chief scientist for AI at Google Cloud, emailed colleagues. "Weaponized AI is probably one of the most sensitized topics of AI—if not THE most. This is red meat to the media to find all ways to damage Google." She was right. When news broke of the contract, uproar ensued. "Don't be evil" had

once been Google's motto, but by the time of Project Maven it had been replaced by a more herbivorous pledge to "Do the right thing." Nevertheless, many of the company's more than eighty thousand workers were shocked that Google was involved in the military's lethal drone assassination program. Around four thousand of them signed a petition demanding "a clear policy stating that neither Google nor its contractors will ever build warfare technology." A handful resigned.

Amid the furor, no one appears to have noticed that Project Maven fit into the grand tradition of many other high-tech weapons projects: ecstatic claims of prowess coupled with a disregard for real-world experience. The "full motion video" to be processed through Google's technology was to be provided by Gorgon Stare, a system of pods arrayed with cameras mounted on Reaper drones. "We can see everything," the *Washington Post* had announced breathlessly in a report on the system's alleged capabilities back when it was first unveiled. But, as the Air Force discovered, that was not true. A testing unit at Eglin Air Force Base revealed in a 2011 report that, among numerous other deficiencies, the cameras could not "readily find and identify targets," and its transmission rate was too slow. The testers concluded that it was "not operationally effective" and should not be deployed in Afghanistan; the Air Force sent it anyway. "A lot of money has gone into it, and I'm telling you right now the fielded stuff still can't do it," a former fighter pilot and longtime Pentagon analyst familiar with the ongoing program informed me. After years of expensive effort, the "ATR [automatic target recognition] just doesn't work. Even driverless cars," he emphasized "are relatively easy problems to solve compared to ATR." Another combat veteran, now with a Pentagon agency working on these issues, told me that the AI developers he works with didn't seem to understand some of the requirements for the technology's military application. "I don't know if AI, or the sensors that feed it for that matter, will ever be capable of spontaneity or recognizing spontaneity," he said. He cited a DARPA experiment in which a squad of Marines defeated an AI-governed robot that had been trained to detect them simply by altering their physical profiles. Two walked inside a large cardboard box. Others somersaulted. One wore the branches of a fir tree. All were able to approach over open ground and touch the robot without detection.

Despite the moral objections of Google employees, Project Maven did not die. Following the company's withdrawal, subcontracts were picked up by Microsoft, Amazon, and Palantir, among others. There was no public announcement, but the new deals were unearthed by a Google employee who had quit in revulsion. Jack Poulson, a former Stanford professor who specializes in advanced supercomputers, had left academic life in 2016 "to go where the actual experts were, which clearly was Google at the time," as he put it to me. Once installed as a senior research scientist, he found that he disliked the corporate culture, despite the informality of the office, where some of his colleagues worked barefoot. It was their "sense of righteousness" that grated him. "Google culture was defined by an exceptionalism that they are the uniquely ethical company," he told me, but he found they were prepared to "look the other way" when the company did "terrible things." Poulson's final decision to quit was spurred by the information that, while eagerly pursuing Pentagon business, Google was simultaneously working with the Chinese government to build Dragonfly, a version of its search engine that would blacklist certain search terms, such as "human rights" and "student protest."

The following year, Poulson was invited to a meeting at a well-known conservative think tank, alongside tech CEOs, high-ranking military officers, and intelligence officials. They were intrigued, he remembers, by what they regarded as Poulson's eccentric attitude toward defense work. One and all, they were eager to promote tech involvement in defense. "It was very shocking to me, the degree to which very senior U.S. military officials were desperate to work in the tech industry," he recalls. "At one point, a high-ranking general was saying that one of the goals was to have at least a hundred flag officers in executive positions by the end of the next year."

Conscious that few outside, or even inside, the business understood the extent of tech's role in the military-industrial complex, Poulson, along with other dissidents, launched a website called Tech Inquiry. "There was so little coverage," Poulson told me. "I started looking into some of these organizations and the contracts with the Defense Innovation Unit—it just seemed like such a zoo of companies." Many of the details were obscure, buried in contracts accessible only through the relentless pursuit of Freedom of Information requests.

Google's retreat on the Maven contract sparked outrage in the burgeoning defense-tech complex. Thiel called Google's stance "treasonous," while Schmidt said he "disagreed" with the decision. *The Atlantic* ran a story headlined THE DIVIDE BETWEEN SILICON VALLEY AND WASHINGTON IS A NATIONAL SECURITY THREAT. Amazon overlord Jeff Bezos, during a live event with *Wired,* huffed that "if big tech companies are going to turn their back on the U.S. Department of Defense, this country is going to be in trouble."

Bezos had a personal interest in the issue. Amazon had accelerated tech's intrusion into the defense complex back in 2013, winning a $600 million ten-year contract from the CIA for uses of the Amazon Web Services cloud originally built to host commercial-customer transactions. Cloud computing was becoming the new profit frontier for corporations such as Amazon, Microsoft, Oracle, and Google, which had sloughed off the moral qualms of the Maven episode to bid for a slice of the $9 billion Joint Warfighting Cloud Capability contract. Though they still reaped vast profits from commercial services, corporate eyes were turning to the government as a stable source of bountiful revenue. If indeed there ever had been a tech-Pentagon divide, it was disappearing fast.

Much has been made of the impact of drones on the war in Ukraine, the implications being that they represent a revolutionary advance in weapons technology. But arguably the most common and effective drones used by both sides have been simple, cheap devices whose parts are available to any consumer with an Amazon account, and jerry-rigged to carry small bombs or shells. They resemble the homemade IEDs deployed with deadly effect by insurgents in Iraq or Afghanistan more than high-tech DARPA products.

Nevertheless, the tech industry has been eager to show that it can do better. This past December, Anduril, another tech enterprise financially godfathered by Thiel to pursue the national security market, announced that it had developed Roadrunner, a small jet-powered drone purportedly capable of autonomous detection and destruction of jet-powered threats, including drones, before returning to its base, costing "in the low hundreds of thousands," according to *Newsweek*. The announcement included a slick video of a Roadrunner flawlessly completing a test. As is customary, the media applauded "America's latest game-changing weapon." Veteran Pentagon analyst Franklin "Chuck" Spinney had a less starry-eyed reaction. "Marketing drivel," he told me, expressing curiosity about how the drone was depicted landing tail-first in the desert with its exhausts hardly kicking up a plume of dust. (Anduril has not given any explanation for the apparent discrepancy.) Meanwhile, the Pentagon insider who alerted me to the cardboard box maneuver, and who has had the opportunity to review Anduril's products, described them as "the F-35 of that world . . . complicated to operate and too pricey." Undeterred, the Special Operations Command awarded Anduril a billion-dollar contract for a counter-drone system in January 2022.

Anduril, named, like Palantir, in a whimsical reference to *The Lord of the Rings*, is the brainchild of Palmer Luckey. As a teenager, Luckey developed Oculus, a virtual-reality headset company that he sold to Facebook for $3 billion. He then focused on defense work, lamenting that people with the relevant tech skills to build the weapons of the future were "largely refusing to work with the defense sector." After selling a border-surveillance system, Luckey turned his attention to drones and other weapons systems, all while harboring more ambitious visions for the future of warfare. Speaking to a sympathetic CNN interviewer in 2018, he revealed that Anduril was working on an "AI-powered sensor fusion platform" to "build a perfect 3D model of everything that's going on in a large area." Soldiers in the future, he predicted confidently, would be "superheroes'' with "the power of perfect omniscience over their area of operations, where they know where every single enemy is, where every friend is, where every asset is."

However fanciful the idea might seem, Army leadership was sufficiently seduced by the prospect of perfectly omniscient soldiers to award Microsoft a $21.9 billion deal in 2021 for an Integrated Visual Augmentation System (IVAS) based on the HoloLens headset. Unfortunately, a report from the Pentagon's director of operational test and evaluation released in 2023 revealed that the majority of soldiers testing the system "reported at least one symptom of physical impairment to include disorientation, dizziness, eyestrain, headaches, motion sickness and nausea, neck strain and tunnel vision." They also complained of the system's

> poor low-light performance, display quality, cumbersomeness, poor reliability, inability to distinguish friend from foe, difficulty shooting, physical impairments and limited peripheral vision.

Congress withheld most of the procurement funding, but the Army promptly handed Microsoft another $125 million to tweak the system, with the price tag climbing to just under $23 billion. (The Army claims that "soldier feedback was positive" in its in-house tests on the tweaked system.)

The rewards of IVAS, at an estimated $60,000 per soldier, are no doubt welcome to Microsoft, given that the market for the HoloLens has been faring badly, but the giant corporation has a far more glittering prize in sight since the ecstatic reaction to OpenAI's unveiling of ChatGPT. Not everyone has been impressed. Commenting on the alleged threat of super-intelligent machines, Meta's chief AI scientist Yann LeCun suggested that such alarms were "very premature until we have a design for a system that can even rival a cat in terms of learning capabilities, which we don't have at the

moment." Nate Koppikar, who has made hundreds of millions of dollars betting against tech as co-founder of the investment firm Orso Partners, is another skeptic. "What's been going on in tech since 2016 has just been, for lack of a better term, a bunch of bullshit," he told me. He pointed out that the sector has been losing money and cutting staff following a bubble that finally burst in late 2021, losing industry investors roughly $7.4 trillion. "And then all of a sudden we're sold this promise of AI as the next big thing." He believes that the fears over evil AI robots "like Arnold Schwarzenegger's Terminator" are part of a marketing campaign to convince us of the awesome power of the technology which, as he pointed out, suffers from inherent defects, including a propensity to make things up, a seemingly intractable tendency known in the industry as "hallucinations." He pointed to Palantir, which, he said, had been losing money prior to going public in late 2020. (The stock went from a high of $45 in January 2021 to just under $8 two years later.) "So they've completely flipped the script," he said. "This year they turned themselves into an AI company."

I was curious about Palantir, whose stock indeed soared amid the 2023 AI frenzy. I had been told that the Israeli security sector's AI systems might rely on Palantir's technology. Furthermore, Shin Bet's humiliating failure to predict the Hamas assault had not blunted the Israeli Defense Force's appetite for the technology; the unceasing rain of bombs upon densely packed Gaza neighborhoods, according to a well-sourced report by Israeli reporter Yuval Abraham in *+972 Magazine,* was in fact partly controlled by an AI target-creation platform called the Gospel. The Gospel produces automatic recommendations for where to strike based on what the technology identifies as being connected with Hamas, such as the private home of a suspected rank-and-file member of the organization. It also calculates how many civilians, including women and children, would die in the process—which, as of this writing, amounted to at least twenty-two thousand people, some 70 percent of them women and children. One of Abraham's intelligence sources termed the technology a "mass assassination factory." Despite the high-tech gloss on the massacre, the result has been no different than the slaughter inflicted, with comparatively more primitive means, against Dresden and Tokyo during World War II.

To determine whether Palantir, which CEO Alex Karp has proudly proclaimed "stands with Israel," was playing a role in the mass killing, I contacted the company, but received no response. So I turned to AI. I first asked OpenAI's ChatGPT, which told me that it had no information because its training had ended in January 2022. I then asked Google's AI platform, Bard, which confirmed that there was indeed such an arrangement. "In 2019, Palantir announced a new partnership with the IDF to develop AI-powered tools for the IDF," it replied. "These tools are designed to help the IDF to identify and track potential threats, and to make better decisions about how to allocate resources." When I asked for the original announcement, it sent me a professional-looking press release, complete with supportive quotes from Karp and the IDF's chief of staff—exactly the information I sought. But there was a problem: it was a hallucination. (It was mistaken about the IDF chief of staff.) No such press release had ever been issued.

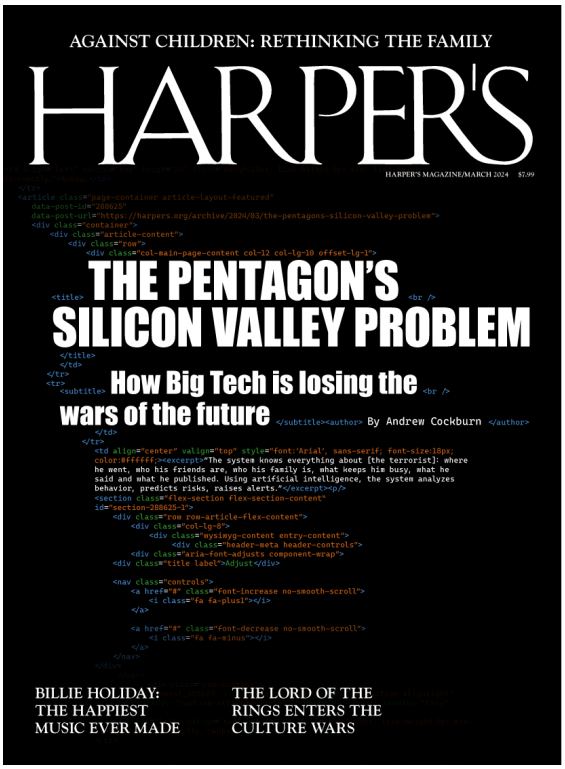

From the