ANNALS OF TECHNOLOGY

# COMPUTER CHIPS ARE THE MOST MAGNIFICENT THINGS HUMANITY HAS EVER MADE

*They power our phones, cars, dishwashers—and the A.I. systems that prop up the economy. They're also at the heart of fierce geopolitical rivalry.*

By Gideon Lewis-Kraus

October 19, 2026




"We don't have a rating scale for complexity, but rockets are maybe a five," an investor said. His partner added, "Semiconductors are an eleven." Photograph by Aleksey Kondratyev for The New Yorker

Save this story

**Listen** • 75 minutes

There are a handful of truly monumental collective achievements in our history: Machu Picchu, the Great Wall of China, the Pyramids of Giza, and perhaps, if we're being generous, the Roman aqueducts. Modern global civilization has made one

contribution to this eternal pantheon. It is not the moon landing. It is the semiconductor supply chain.

This strikes almost all of us as ridiculous. A microchip is a bauble of metal and refined sand. It doesn't appear to have epic proportions. Upon close examination, however, the epic proportions are there in miniature. Imagine a perfect replica of a sprawling metropolis, a Tokyo or a São Paulo, engraved at atomic dimensions on a tile the size of a fingernail. The urban details would be faithfully rendered to the last golf ball or egg—each reduced to approximately the bulk of a hemoglobin molecule, or the distance grown by a human fingernail every five seconds. If a chip were merely a sculpture, this chiselling would represent a considerable achievement. But a chip is, in fact, a functional city for a population of busy electrons. They buzz about, flying our airplanes, powering the stock market, and, increasingly, thinking on our behalf. Unlike rockets, which are produced in relatively meagre quantities, microchips are stamped out by the billion. This is the most impressive thing about them. Or, maybe, about us.

Microchips are mass-produced, like any tchotchke, at factories. Most of these factories, and the best of them, were once situated in the United States. By the nineteen-eighties, that began to change. American semiconductor firms were threatened by the emergence of Japanese rivals. Our domestic industry petitioned the government to intervene. Economists shrugged. The consensus at the time was in favor of free trade. If a good could be inexpensively produced elsewhere, everyone benefitted. American consumers deserved bargains on digital watches, pocket calculators, and primitive computers. The most legendary, if apocryphal, expression of this attitude was credited to a George H. W. Bush Administration factotum: "Potato chips, computer chips, what's the difference? They're all chips."




*"My parents are gonna love you as soon as they calm down from hearing the doorbell."*

Cartoon by Asher Perlman

Open cartoon gallery

This wasn't entirely a joke. A columnist for the Los Angeles *Times* pointed out that the largest domestic chipmaker, Intel, was "roughly the same size as PepsiCo's Frito-Lay Inc. subsidiary, the largest potato chip maker." (An executive at Utz got defensive, complaining that "there's quite a bit of technology in modern potato chip manufacturing.") Widgets were widgets, was the point. If we didn't care about where salt-and-vinegar chips came from, why should we care about their silicon-flavored cousins?

Today, we very much care. Chips are the foundation of A.I., among other things, and A.I. is propping up our entire economy. The most valuable company in the world is the chip designer Nvidia. Last year, the company's stock single-handedly accounted for about fifteen cents of every dollar returned by the S. & P. 500. From an economic perspective, the problem with chips nowadays is that there are never enough. A car might contain two thousand of them. When the pandemic caused supply-chain issues, automotive manufacturers found that they'd neglected to reserve a sufficient number. The government pressured chipmakers to reprioritize order fulfillment, safeguarding the birthright of Americans to buy new trucks. The chips that make Claude and ChatGPT talk are more specialized than those in a Chevy, and the chip industry has responded by reallocating its resources to these product lines. If basic chips grow scarce, it will not be long before the average homeowner cannot afford a dishwasher, let alone a phone. If the A.I. backlash seems fervent now, just wait until the dishes pile up.

But the financial interests pale in comparison with the geopolitical ones. The vast preponderance of the world's most advanced chips are made in Taiwan, a tropical island best known in America for its bubble tea, the Din Tai Fung restaurant chain, and its tense proximity to a superpower that regards it not as an independent nation but as a temporarily lost belonging. The prospect of a Chinese invasion has long loomed over Taiwan's horizon.

If China were to invade, or even throw up a naval blockade, life as we know it would continue for about the time it takes to manufacture a chip—a season or two. Then the inventory would thin out. The sticker price of an iPhone might be confused with that of a luxury watch. Not that anyone would care. Stocks of Wi-Fi routers, insulin pumps, pacemakers, and credit-card readers, along with cars and dishwashers, would dwindle and perhaps vanish. Tech equities—for all intents and purposes, the stock market itself—wouldn't be worth the paper they're printed on. Bloomberg Economics has estimated that just the first year of a war over Taiwan could cost the world economy $10.6 trillion, or roughly a tenth of global G.D.P. Pat Gelsinger, a former C.E.O. of Intel, said to me, "You think the closure of the Strait of Hormuz was

frightening? That was trivial. It's twenty per cent of the oil. Taiwan is eighty per cent of the supply chain of the world. A three-week brownout there would have a larger economic impact than the Great Depression."



The Dutch firm ASML, headquartered in Veldhoven, is Europe's most valuable company.

Photograph by Alastair Philip Wiper for The New Yorker

At stake, in a longer term that might not be so long, is the global balance of power. Chinese chipmakers remain about five years, if not more, behind their Western counterparts. Should American A.I. companies find themselves restricted to alternative outputs—to the small fraction of suitable chips produced in Korea or at home—our current lead in the arms race would almost instantly totter. It might not be the worst thing if our college students resumed writing their own essays. But a lack of chips might also spell the end of our military advantage.

As a result, Washington has recently come around to the idea that the provenance of semiconductors matters more than the provenance of Pringles. A nostalgic consensus has emerged in favor of their domestic production. To a casual observer, the construction of a chip factory might seem straightforward: you buy some land and throw some money at a contractor to build a huge room to house some machines. It is, alas, not so simple.

Taiwan Semiconductor Manufacturing Company, or T.S.M.C., makes almost all of the world's advanced chips, and has a market capitalization of around two trillion dollars. Of the ten most valuable companies globally, it is one of only two that are not American tech giants. (The other is Saudi Aramco.) Unlike T.S.M.C.'s peers—household names such as Apple, Google, and Amazon—the company has never cultivated much of a public-facing brand identity. Its fabs, as the factories are called, are only begrudgingly opened to clients, and even they are afforded a look only at their own production lines. The extreme reticence is by design. Some of the chip designs that the company executes, like those which animate an iPhone or a Mac, are inestimably precious trade secrets. Others, like the patterns for Nvidia's top of the line, are treated like national-security assets, and are subject to American export controls.

T.S.M.C. is headquartered in the Hsinchu Science Park, a manicured enclave on the outskirts of a coastal city about an hour from Taipei. Its hulking fabs are neither hardened nor hidden. Some of them have panoramic views of white windmills, which dot the coastline of the Taiwan Strait like slender sentries. China lies a hundred miles

or so over the water. The corporate nerve center is clad in the company palette of muted grays and bright-red accents. When I first approached the lobby, just before seven o'clock one humid morning in June, a maintenance man was showcasing the firm's attention to detail. He stood knee-high in a decorative pond, carefully pushing a waterlogged broom along its floor to sweep away invisible dirt from between elevated plinths bedecked in artificial roses. The interior décor consisted in its entirety of an oil portrait of Morris Chang, the company's founder. He appeared with his signature pipe, in an untidy study surrounded by piles of books and lit by a luminous silicon wafer.

A senior receptionist stood stern and unblinking behind an enormous desk, presiding over narrow entry turnstiles. Various items are prohibited across the complex, including phones, smartwatches, recording devices, and, curiously, blank double-sided pieces of printer paper. Given that I had yet to proceed beyond the lobby, I ventured to ask permission to photograph the painting. The answer was an immediate, polite, and decisive no.

The original computers were essentially thickets of light bulbs. The basic logic of computing is that practically any set of instructions can be executed by an interconnected constellation of switches that can be turned on and off. Those first computers, developed during the Second World War, evoked the basement warehouses of a chandelier wholesaler. They required a platoon of soldiers to race around replacing bulbs as they burned out. In theory, these devices could accomplish tasks of great complexity. In practice, even a relatively basic job might result in a cairn of charred copper and broken glass.

After the war, at Bell Labs, the physicist William Shockley and two colleagues invented a switch that was smaller, faster, and didn't explode. The transistor, as it was called, exploited the peculiar disposition of a kind of element known as a semiconductor: it could be flipped on to conduct electricity, the way a metal would, and it could be turned off to block the current, like glass. The little gizmo earned its inventors a Nobel Prize and almost immediately improved lives. The portable transistor radio, which appeared just in time for Christmas in 1954, became a consumer blockbuster. The computers of the future would no longer need to be enormous things that caught on fire. Their potential seemed limitless.

There was, however, a devilish obstacle. Computing relies on circuits, or transistors that have been wired together with other electrical components. The cleverness of a circuit depends on two factors: how many transistors it contains and how densely they are interconnected. Transistors, unlike light bulbs, could be added to circuit designs in promiscuous quantities. Unfortunately, they still needed to be soldered together by hand. This process took time and money and tended to be flaky. One little miswiring could ruin everything. As circuits got larger, the difficulty of the interconnective task increased exponentially. This was called the "tyranny of numbers," and it was demoralizing. By the nineteen-fifties, we knew that getting to the moon would require computer guidance. A circuit that could do it, with ten million components, could be sketched on paper. But no one could actually build it.

In the winter of 1958, two men independently figured out how these paper circuits could be made flesh. The old method assembled circuits from individual transistors, as if a jumble of disconnected body parts were stitched together. The new method took a monolithic piece of semiconducting material, as if it were a block of marble, and carved a figure out of it. Had the initial choice for substrate prevailed, the patch of cherry orchards around Palo Alto would be called Germanium Valley.

The origin story of Silicon Valley contains two contradictions. The first is about where it happened. One of its godfathers, Jack Kilby, was based not in California but in Dallas, at Texas Instruments. Silicon Valley was an idea, not a place. The second relates to the role of competition in innovation. The other godfather, Robert Noyce,

began his career under Shockley, the transistor pioneer (and famed scientific racist). Fed up with Shockley's managerial despotism, Noyce—along with seven colleagues, who became known as the Traitorous Eight—defected to found a rival firm, Fairchild Semiconductor, and then defected again to found Intel. In the ensuing decade, lawyers at Texas Instruments and at Fairchild fought a bitter court battle over the patent. Subsequent generations of founders drew from these tales the moral lesson that every great company originates in betrayal and is sustained by combat. In fact, this wasn't at all how Kilby and Noyce felt. They cared more about solutions than about heroics, and were delighted to share the credit.



Employees at a factory that makes ASML's extreme-ultraviolet machine. "Internally, there's a pride in the fact that you have no idea how complicated what we do is," an engineering executive at the company said.

Photograph by Aleksey Kondratyev for The New Yorker

The first integrated circuits on the market were compact, trustworthy, and completely unaffordable. The only plausible customer was the one organization with the desire and the budget to pursue an irrationally difficult goal at any cost: the federal government. For the nascent chipmakers, the military-industrial complex was a godsend. Aerospace and defense contracts paid for the capital-intensive installation of a manufacturing base. But executives chafed at the limitations of Washington's monopsony: their engineers didn't like answering to bureaucratic brass; sectors overreliant on government largesse get flabby; and only widespread demand would

allow them to achieve economies of scale. The executives courted commercial interest, and this order volume translated into improved efficiency, accelerated research, and lowered prices. In 1971, Texas Instruments released the Pocketronic calculator. It offered four cutting-edge features: addition, subtraction, multiplication, and division. By the end of the decade, according to T. R. Reid's "The Chip," the United States had more calculators than people.

In 1964, someone asked Noyce's colleague Gordon Moore to predict the technology's trajectory. He'd noticed that, in each of the previous few years, transistor density—the number that could be shoehorned into the same footprint—had doubled. Moore casually reckoned that this trend would continue. At the time, a state-of-the-art chip had sixty transistors. By his logic, the chip of 1975 would accommodate sixty-five thousand—a ludicrous estimate. When the bet was settled, the newest chips boasted sixty-five thousand five hundred and thirty-six. If the automobile industry had innovated at the same pace, Moore calculated a few years later, "we would cruise comfortably in our cars at 100,000 mph, getting 50,000 miles per gallon of gasoline." The wildest part is that the upgrades were essentially free. Moore went on, "We would find it cheaper to throw away our Rolls-Royce and replace it than to park it downtown for the evening."

The spell hasn't been broken. The phone in your pocket is several orders of magnitude more sophisticated than the computers used for the Apollo mission. One AirPod could easily calculate a path to the moon. A car-key fob would suffice. An electric toothbrush. There are Rolls-Royces that we do, in fact, routinely throw away. The Apollo mission's chips were less powerful than those currently inside a digital pregnancy test. We squander moon-landing capability to display a plus sign.

The Dutch city of Eindhoven was founded on the mastery of artificial light. In the late nineteenth century, two brothers, Gerard and Anton Philips, started an incandescent-bulb outfit that would eventually become one of the largest electronics conglomerates in the world. The first jerseys of the local soccer team featured a light-bulb emblem, and graffiti light bulbs are still spray-painted on the city's underpass walls. Nearby, the leafy suburb of Veldhoven is now home to ASML, Europe's most valuable company. What the firm sells is rarefied light.

ASML, originally named Advanced Semiconductor Materials Lithography, is perhaps best known for not ringing a bell. In 2024, the *Wall Street Journal* called it "the most vital company that you've never heard of." The firm wears its obscurity as a badge of honor. Alex Schafgans, an engineering executive at ASML, told me, "Internally, there's a pride in the fact that you have no idea how complicated what we do is. If someone at a barbecue asks me to explain my job, it takes me eight minutes." A finished microchip from T.S.M.C. is the culmination of thousands of steps. Each is indispensable in its own way, but the most indispensable is the chip's passage through an ASML machine.

Moore's Law was never an actual law, like gravity, but an observed regularity. Chips improved on a regimented schedule. This observation became a promise: every two years, the industry would deliver a product that was twice as good, at no extra cost. The rest of the world came to count on this. Apple, for instance, takes for granted that it can design a new phone using powerful chips that do not yet exist. These fictional chips are then summoned into existence by the delusional commitment of more than two million chip-industry employees to the perpetuation of a self-fulfilling prophecy.

A chip's transistors, for example, were once interconnected with microscopic lines of aluminum wire. As transistors got smaller, the channels between them became proportionally narrower, and aluminum's performance hit a wall. Copper would have been a superior choice, but it leached into silicon and ruined chips; its very presence in fabs was strictly verboten. But I.B.M. researchers, inspired by the metalwork of

medieval Damascene artisans, invented a technique to protect silicon from copper. Moore's Law was given a reprieve.

The fundamental insight behind the mass production of semiconductors is that a chip can be printed, like a photograph, with light. A chip's architectural blueprints are traced on a stencil, or "mask." When light is projected through the stencil, like sun through a stained-glass window, the chip's ghostly outline is inscribed on a slab of silicon that has been coated with a light-sensitive chemical. Gelsinger, the former Intel C.E.O., resorted to theological comparison: "At the center of a fab is lithography, and at the center of lithography is light, and God said, 'Let there be light.' "

At the dawn of the industry, this procedure could be cobbled together from normal commodities. Noyce, of Fairchild, went to a photography store and found a few serviceable lenses in a crate of antique movie cameras. The stencil could be outlined on a pane of conventional glass. The light source was a regular lamp. As transistors shrank and chips grew denser, the process called for greater precision. Chips could be destroyed by even a speck of dust, so factories became impregnable clean rooms, their workers clad in bunny suits. Janky store-bought components were replaced with custom materials and tools. Light was no exception. Lamplight was too coarse an instrument for fine patterns. At shorter wavelengths, however, the light could be wielded like a scalpel.

The industry marched its way down the spectrum to violet light, the shortest wavelength we can see, and then off the cliff of visibility. Beyond that was "deep ultraviolet," where our reliance on God-given light gave way to excimer lasers made with krypton or argon-fluoride gases. This solved one problem but created others.

The transient air pressure of a passing thunderstorm, or the excessive production of cow patties, spoiled the chips. These issues could be patched, but the underlying threat was existential. By the early nineteen-nineties, deep UV seemed like reality's best and final offer. We were poised to carve greater detail into silicon until approximately 2005. The chips we had that year were going to be the best we'd ever have. In that event, we probably wouldn't have the iPhone, not to mention Uber or Spotify or Netflix, which perhaps would've been for the best. But we'd also likely lack reliable five-day weather forecasts, automatic emergency braking, pocket ultrasound, rapid genetic sequencing for sick babies, and images of black holes.

The bulk of contemporary life rests on a decision to bet the farm on a technology that almost certainly could not be realized. Executives at Intel, with the federal government's help, poured money into "extreme" ultraviolet light. There were about a million reasons it couldn't be harnessed, among them the fact that the nearest natural source of EUV is the sun. Owing to both technical and political concerns—Japanese lithographers were off the table—the only hope was offered by ASML, then a relatively obscure shop. Jan Mulkens, a physicist, joined the company during this tumultuous period. He'd left his first job interview, in 1992, thinking, *These guys are completely crazy. I must work here*. ASML runs on the belief that if you recruit people who don't know enough to recognize that the job before them is hopeless, they might surprise you.

No single institution, no matter how crazy its employees, could do it alone. Even if you could somehow conjure terrestrial starlight on an industrial scale, it would immediately disappear into anything—a stray particle of air, say, or the glass lenses used to collect light. ASML collaborated with the storied German lens crafter Zeiss.

Mulkens, who is also an amateur cyclist, despite having broken his shoulder four times, said, "It is five hundred and ninety-two kilometres to Zeiss, and I have done it more than one hundred times." Jos Benschop, a sly, wolfish vice-president at ASML, assumed a mask of theatrical German intensity: "When you design things together, you need to see each other, you have to have almost daily contact—you need to look into each other's eyes." The collaboration produced mirrors that are the smoothest man-made objects in the known universe. If you blew one up to the size of the planet, the highest mountain would have the elevation of a face-down playing card. Other mirrors that Zeiss contributed could direct a beam of light to hit a dime on the moon; in fact, they could successfully target either the left or the right half of that dime.



An ASML employee at a company facility in San Diego. Photograph by Aleksey Kondratyev for The New Yorker



A closeup of an ASML machine. Photograph by Aleksey Kondratyev for The New Yorker

In 2006, Samsung placed the first order for an EUV machine, but warned that the cost was worth it only if a certain number of chips could be printed per hour. A viable

prototype needed to produce a minimum of a hundred watts of light. At the time, ASML was stalled out at five. The company blew the deadline by seven years. Samsung, understandably, began running out of patience.

By the finally final deadline, Schafgans, then a researcher, had been "fighting all these ghosts in the machine for years." One day, at his wits' end, he groped inside the machine for a little gauge that had been installed to control a particular setting. It wasn't the sort of crank that was meant to be turned, but Schafgans had the impulse to monkey with it. He said, "Imagine me, in a clean-room bunny suit, holding my laptop with one hand, with the other reaching into that massive machine with a little screwdriver to adjust a dial that was never even meant to be adjusted." His instinct was to turn it in what was, mathematically speaking, the obviously incorrect direction. For some reason that remains largely incomprehensible, it worked. Schafgans got a big promotion.

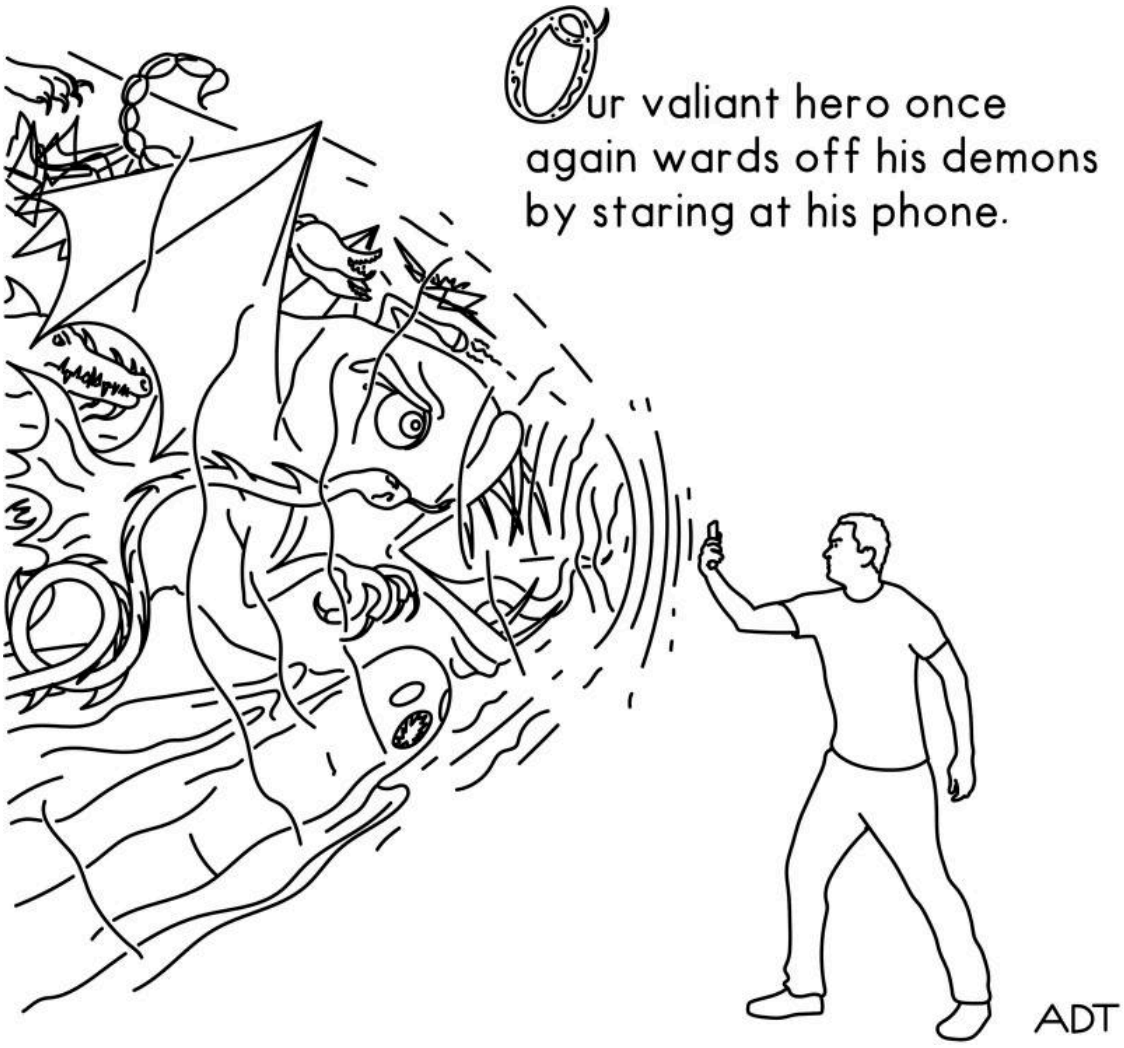
Three years later, in 2019, a package from Samsung arrived in Veldhoven. To the untrained eye, it was just a phone, but industry veterans saw something inconceivable. A chip in the Apollo mission's guidance computer had six transistors, each roughly a third the width of a human hair. Inside the phone was the first EUV-printed chip, which had more than ten billion transistors, each of them narrower than a virus. When ASML was founded, one gigabyte of memory cost almost a million bucks. A gig of EUV-printed memory costs only a couple of dollars. Mulkens spoke as if these dizzying statistics were mere by-products of his colleagues' valiant struggle against the laws of physics: "Moore's Law, I do not give a shit about it. It's good for marketing. People who think about it never worked on it themselves."

ASML's main plant is dominated by a forbidding tower of slate gray which looks like a giant chip, a resemblance that's perhaps a bit on the nose. Viewed from an observation deck, the campus is constrained between major roadways, with agricultural fields and gabled roofs beyond. At the company store downstairs, the most popular item is a Lego set of the EUV machine. The passion project of one employee, it has almost a thousand parts and costs two hundred dollars. The actual machine has millions of parts and costs up to four hundred million dollars.

No single person can claim to comprehend an EUV machine in its entirety. The apparatus is roughly the size of two stacked fire trucks and is delivered to customers, such as T.S.M.C., in a fleet of seven 747s. When installed, it's sheathed in cabinetry. A guide and I badged into one of a few dozen restricted cabins where the machine is assembled—an endeavor that takes about twenty weeks of round-the-clock work. The one before us was almost finished. Undressed, it had the appearance of a biomechanical hallucination: its monstrous exterior was trussed with a tangle of silver wires, cables, tubes, and dials, as if Optimus Prime had been put on life support. From a raised platform on the back side, a supervisor pointed to a small opening where some T.S.M.C. technician will begin the chip-printing process by feeding the machine an ingot of tin.

Tiny droplets of the molten tin are fired out of a cannon into a vacuum chamber, where they are shot—sixty thousand times a second—by unerring lasers. When the tin fog is ignited, the explosion produces a plasma fifty times hotter than the surface of the sun. The process, in other words, commences with the industrial production of miniature supernovas. These supernovas emit EUV that must be collected by Zeiss's mirrors, which cannot function if their surface is marred by even a nanometre of tin debris. The only way to keep the mirrors pristine is to pump a continuous hurricane of hydrogen gas into the supernova's blast radius. The gathered light ricochets through a hall of mirrors before it is bounced onto the blueprint of the chip. Making this happen a single time took decades of research. The profitability of ASML's customers relies on it happening near-flawlessly twenty-four hours a day. Recurrent shutdowns, it's not much of an exaggeration to say, would derail the global economy.

A few days before I arrived in Veldhoven, an internal event featured a conversation between ASML's C.E.O., Christophe Fouquet, and Elon Musk. In the past, Musk has threatened to bypass any "rate limiter"—a "bottleneck" or "choke point" in chip production—with his own ingenuity. ASML is the sole purveyor of this particular tool, without which there are no advanced chips, and Musk was presumably referring to the fact that the company puts out only about fifty machines each year. He seems to regard the machines as overgrown photocopiers, and doesn't understand why the company doesn't make more faster. Musk isn't the only kibbitzer who believes that a European firm, with its soft Continental indulgence of extended summer holidays, has taken the future hostage. This summer, after the embattled A.I.-finance guru Leopold Aschenbrenner's hedge fund almost collapsed, he bravely coped by making a four-hundred-million-dollar investment in a U.S.-based lithography competitor. Frustrated American politicians, according to the Dutch journalist Marc Hijink, refer to ASML as "*that company*."




Cartoon by Adam Douglas Thompson

Open cartoon gallery

The reason that nobody, in America or elsewhere, can "just" build an equivalent machine is that there's no "just" about any of it. Some components are novel inventions. Others are transcendent renditions of existing products. A basic metal valve, for example, is available at any hardware store. The valves in ASML's "droplet generator," however, need to be able to handle molten metal under literally oceanic pressures, and to do so day and night for a year without wilting. The company had to find a supplier that could make such a valve. But venders didn't have a financial incentive to do so, since ASML wouldn't be buying the widget in normal widget quantities. ASML's ability to persuade a vender anyway relied on an extended relationship, with accumulated reserves of trust on both sides. The basic model of

competitive markets which one finds in an introductory economics textbook—where independent firms engage in cold arm's-length transactions motivated by unadulterated self-interest—does not apply to the industry. It better resembles a kind of coöperative, or a planned economy.

There is a famous story about an ASML engineer who explained to his Chinese rivals that they couldn't build an EUV machine even if they were handed the blueprints. Marc Assinck, a trim coffee obsessive with oversized Ray-Bans who has played multiple roles at ASML, told me that the very concept of such a blueprint is a fallacy: "There simply is no 'file' of our machines. The intellectual property is spread out across the entire system." He stopped in front of a glass machine component with recondite thermodynamic properties. "Zeiss doesn't have this recipe," he said. "Their glass supplier does." Many of these parts have lead times of two years.

If many aspects of the industry can't be matter-of-factly cloned, it's because they're not particularly industrial. Some of them are properly artisanal. Assinck said, "It relies on machinists, mechanics, and other craftspeople with unwritten knowledge." He pointed to a recessed thingamabob supplied by a family shop nearby. Hundreds of pinprick holes had been bored into the metal undercarriage, and each one had to be deburred by hand.

A major module of the EUV machine is fashioned at an ASML facility in Connecticut that has the aura of an alchemist's atelier. When I visited, in July, arcane prisms rotated on medieval-looking polishing machines. The facility hires the sorts of people who fix their own motorcycles. One young employee sat at a workbench in front of a block of milled amber glass that appeared extraterrestrial. The block, which bears a chip's blueprint, levitates inside the machine. It moves with the precision

necessary to trace a strand of DNA, even while enveloped in a robot that accelerates with forces great enough to tear an airliner apart nine times over. The entire apparatus is like an MRI machine that is expected to scan hairline fractures in the middle of a violent plane crash. The engineer proceeded by feel. He paused now and again to tap his fingers gently against a particular spot. It seemed less like engineering than like reflexology. This felt poignantly ironic. When you dig down to the very bedrock of A.I., you discover that underneath all of its cognitive power is a dude drumming his fingers with occult proficiency.

Just before 8 A.M. every day of the year, the gates of the Hsinchu Science Park are breached by a legion of commuters arriving on puttering scooters. T.S.M.C.'s home is lush and manicured, but it feels more like an autonomous zone than like a corporate complex. When I visited, this past summer, a guide pointed down the street to the National Experimental High School, from which his son would be graduating the next morning. We turned from Prosperity Road 1 onto Prosperity Road 2. Expert launderers of clean-room bunny suits careened by in branded wagons. There are administrative pavilions, post offices, banks, clinics, and two universities. Some employees live in dorms on site, and all of them have easy access to pools and fitness centers. You rarely have to walk more than a block to visit a sparkling 7-Eleven. T.S.M.C.'s suppliers—tooling companies such as Applied Materials and Tokyo Electron, as well as software firms such as Cadence and Synopsys—maintain their own local embassies.

The visionary behind this vibrant terrarium is Morris Chang. Born in 1931 in the Chinese city of Ningbo, he was raised in wartime flight—moving from Nanjing to Guangzhou, Hong Kong, Chongqing, Shanghai—and resolved to find peace at a university overseas. He enrolled at Harvard, planning to study Shakespeare, but quickly realized that engineering was a more secure path for Asian immigrants, and transferred to M.I.T. At twenty-seven, he accepted a job at Texas Instruments, just in time for the invention of the integrated circuit. Chang was tasked with the problem of yield, or the percentage of finished transistors that worked. The production line that he was given had a yield of approximately zero. With his pipe in his mouth,

Chang devoted himself to operational excellence. Within a few months, the yield exceeded twenty per cent.

Chang was steadily promoted before he was passed over for the position of chief executive, perhaps for bamboo-ceiling reasons. He was in his fifties, and figured that his career had peaked. In the early nineteen-eighties, though, the Taiwanese government recruited him to oversee a new technology initiative. On an episode of the podcast "Acquired," the host David Rosenthal compared the plan to all the countries that say, "We're going to build the next Silicon Valley in XYZ." This, Rosenthal continued, was "probably not going to work."




ASML employees set up a workstation. Photograph by Aleksey Kondratyev for The New Yorker

The postwar Taiwanese economy churned out low-margin goods, such as umbrellas and toys, for foreign clients. Mattel didn't need domestic Barbie factories if a Taiwanese partner could produce dolls more cheaply. Such factories could also then amortize their fixed costs by making adjacent products for other brands. Chang proposed a semiconductor "foundry," which would do for pricey chips what its predecessors had done for cheap junk. At the time, however, the market didn't exist. Chipmakers had plenty of overseas help—in the book "Chip War," the scholar Chris Miller notes they were globalizing long before the word gained currency—but the industry was also known for vertical integration, and did the hardest parts in-house. "Morris, you've had a lot of good ideas in your time," Gordon Moore told Chang. "This isn't one of them." Jerry Sanders, the C.E.O. of the American chip company A.M.D., had once notoriously declared that "real men have fabs."

Real men, as it happened, are delighted not to have to spend twenty billion dollars on a fab. By the end of T.S.M.C.'s first decade, the company was bringing in a billion

dollars in annual revenue. In 1993, a young Taiwanese American named Jensen Huang had an idea for a specialized graphics chip, or G.P.U., to intensify the video-game experience, and founded Nvidia. It wasn't long before design mishaps drove Nvidia to the brink of bankruptcy. Huang sent a personal letter to Chang, and got a call at his office a week later. (The story, in Stephen Witt's "The Thinking Machine," is that Huang yelled, "Everybody shut the fuck up—I've got Morris Chang on the phone.") T.S.M.C. took a chance on his prototype. The specialized architecture of Nvidia's chips turned out to be ideally suited to the development of A.I. In the meantime, Apple developed a novel chip for mobile devices and eventually asked T.S.M.C. to make it. The result was every generation of iPhone that has been produced since 2014.

T.S.M.C. is often described as the heart of Taiwan's "silicon shield." America's unofficial doctrine on Taiwanese independence allows for "strategic ambiguity"—in other words, we might or might not protect the island in the event of a Chinese invasion—but we remain painfully aware of the domestic stakes. We probably need T.S.M.C. more than we need Boeing. Audrey Tang, a Taiwanese civic technologist and a former government minister, told me, "This work isn't an abstraction to people—it's literally what keeps Taiwan on the map."

T.S.M.C. has come to embody national strength and national purpose. A shy, cerebral engineer named Nicholas Lee translates abstract research into procedures that can survive implementation on the factory floor. A machine might need to be tuned, for example, to compensate for low air pressure from an incoming typhoon. He joined T.S.M.C. four years ago, right out of a master's program in chemical engineering at the preëminent National Taiwan University. When I asked how many of his

classmates had done the same, he laughed and said, "At least sixty per cent." He added, "Everyone thinks there's something wrong with you if you don't accept a T.S.M.C. job." To enlist at T.S.M.C. is to give one's parents the gift of familial honor. The company's lower echelons, which tend to be populated by graduates of vocational or technical schools that serve as T.S.M.C. feeders, have the same social coherence. Although it's a "very intense work environment," the challenges were eased by the fact that, as Lee told me, "we all grew up together." One in forty-six babies in Taiwan is born to a T.S.M.C. employee.

Without Morris Chang, we almost certainly would not be talking about A.I., and we probably wouldn't be talking about Taiwan as a sovereign state. When Chang invariably comes up in conversation at T.S.M.C., he is referred to as Founder Chang, or sometimes simply as Founder. Founder is ninety-five, and a folk hero. He now and then appears at the company's annual sports day, for which T.S.M.C. rents out a stadium. As Miller puts it, "Chang wanted to become the Gutenberg of the digital era. He ended up vastly more powerful."

T.S.M.C., I was told, was willing to take me on a tour of a fab called AP6 on two conditions: that I arrive bearing nothing but a passport, and that I dress in garments without so much as a fleck of metal—no zippers, snaps, chapes, grommets, aglets, or ferrules. (Sweatpants were technically acceptable, but might result in loss of face.) The fabs are kept about a thousand times cleaner than a hospital. Their security infrastructure is more impregnable than that of a bank.

T.S.M.C. is a hierarchical, ceremonial place, and the protocol demanded that we first join the fab director for a photo op. Once the picture was taken—the photo itself appeared to remain confidential—we were rushed along an anonymous corridor to a

9 A.M. handover meeting. Every incoming shift begins with a briefing from the outgoing one. Perhaps two dozen people sat in a situation-room arrangement; what I first took for a mirror was actually a video link to an identical configuration in another room in the same complex. Other screens displayed graphs and tables. Jack Chang, a young section manager with moussed hair and stylishly baggy clothing of advanced fabric, whispered to me, "Here, we monitor the efficiency of each production step." He continued, "At this meeting, we ask, Which step was lower than yesterday? Why? How do we minimize any production gaps?"

T.S.M.C. takes a silicon wafer, a shimmering disc with the diameter of a vinyl record, and proceeds to do approximately a thousand things to it. It is coated with an ultrathin film of glass or metal, glazed with photosensitive chemicals, and then passed through a lithography machine. The chemicals are rinsed off, and the pattern on the exposed film is etched in relief by blasts of plasma before being gassed with other chemicals. The wafer is washed, polished, and inspected for flaws. For the most advanced chips, this gantlet is then repeated about a hundred times on a loop, as if an entire city skyline were being raised one story at a time.

If T.S.M.C.'s venders supply the tools, the company supplies the ruthless optimization. Morris Chang founded the business on his faith in the "learning curve." This is an economic principle—put very simply, it's the idea that as you make more of a product, each one costs less—but it's also a cultural one. At a disciplined organization, every product you make is an opportunity to improve how you make it. If T.S.M.C. manufactures the best chips, it is because it has come to manufacture almost all of them.

This method works only as long as employees exercise constant vigilance—over the machines, themselves, and one another. Jack told me, "The culture of this place is that, when something didn't go well last night, you have to tattle to your boss." The meeting attendees at AP6 had the faces of people accustomed to scrutiny. Overnight, one of the polishing tools had experienced an idle time of two hundred and nine seconds. This wasn't an intolerable number of seconds, but it also wasn't ideal. That

morning, the idle time was down to a hundred and twenty-one seconds. No one, Jack told me, was being yelled at today.

In the gowning area, three dedicated valets zipped and snapped me into a purple bunny suit. "Purple means customer or V.I.P.," Jack said. A fourth person handed me a pen and a bound sheaf of blue paper with a starchy crinkle. The booklet, contrived to shed no particles, said "Clean Note: The Highest Clean Level for Clean-Room Use." Once kitted out, we stepped on metal pads that recorded our individual levels of static electricity—stray bodily ions can invisibly crack a wafer's interior—and then proceeded to enter a sealed air-shower corridor, where tridirectional jets treated us to private tornadoes to remove residual impurities.




*"The doctor says I'll be fine if I can avoid unnecessary stress."*

Cartoon by David Borchart

Open cartoon gallery

AP6, which opened in 2023, was T.S.M.C.'s first fab devoted to the entire process of "advanced packaging," an innovation at the cutting edge of the industry. Advanced packaging interlaces several individual chips to form a composite superchip. The challenge is that chips are carved in the miniature realm where matter itself has the irregularity of wood grain. To be stacked together efficiently, they must be painstakingly paired for infinitesimal compatibility. Jun He, recruited from an aerospace shop to help direct the division, compared the manufacturing flow to running the world's largest dating site. He was promoted to vice-president in 2022, right before A.I. blew the roof off demand for the product and relocated his matchmaking service to Hell. He has had to functionally double T.S.M.C.'s advanced-packaging capacity each year since, and if yields are not nearly perfect his customers risk losing money. If he'd known how intense the industry would become, he'd recently told T.S.M.C.'s C.E.O., "I would not have taken this job."

A finished advanced package is typically about half the size of a deck of cards and can have the sticker price of a fancy car. Amy Tsai, a junior manager on the operations team, told me that the total clean-room space—with enough floor area to simultaneously host twenty-two immaculate soccer games—is used for "only our top five customers." She would never name these customers, but there aren't many candidates; the shortlist includes Apple, Nvidia, and Broadcom, which designs A.I. chips for Google and Meta.

We emerged into a long corridor bathed in a sulfurous yellow glow. The floor panels were perforated with a dot matrix of vents that regularly replace all the air in the clean room. Above our heads, where twenty miles of rails had been installed for automated wafer transport, was the orderly traffic of inhuman motion. Robotic claws

whispered along branching tracks, clutching boxy carriers beneath them like prizes in an arcade game. Inside the carriers were the prized wafers. This system, Jack said, had replaced an old method in which a person pushed the wafers in a cart. The human work was less stressful now, a colleague added. He'd once visited a fab with a sign that read "If you bump this cart, your pay will be docked X percent." The facility was so important, one employee later mentioned, that "we pray to the local gods for auspicious function."

Inside the halls, the overhead claws descended to insert wafers into one end of the machines and withdraw them from the other. The only visible evidence of a human hand was paper coverings, which had been taped on the machines to obscure their respective logos. Apparently, there are people who can eyeball an Nvidia Vera Rubin or an Apple M6 from the positioning of, say, a Canon FPA-5520iV LF in relation to a Surfscan SP5$^{XP}$, and T.S.M.C. had mistaken me for one of them. The machines themselves were perhaps less important than the choreography of their administration. As Jun He, the vice-president, put it, "Our key comparative advantage is the management of complexity."

The next day, my guides and I visited AP7, a partially finished facility an hour or two south. The tropical sun was oppressive, but company representatives were prepared with orange parasols, each with a dedicated barcode. Despite slight delays—the excavation of the site had turned up some priceless Neolithic pottery, and the company also had to establish a protected sanctuary for the ring-necked pheasant—the first phase had risen out of an old sugarcane field in only eighteen months. A wiry, blunt thirty-one-year veteran deputy director named K. H. Cheng had come in with a team of shock troops to supervise.



An employee puts on a bunny suit to enter an ASML clean room. Photograph by Aleksey Kondratyev for The New Yorker

The second phase would come online in four to six months. At the loading dock, deliveries arrived in trucks lined up in a queue that had been planned a year in advance. The tools, Cheng explained, were accompanied by vender emissaries who would stay to help for the year it took for T.S.M.C.'s technicians to understand them as well as, or better than, their makers. When I asked if this practice held for ASML's machines, he sighed. "Those people are here all the time, because only ASML's designers know everything," he said. He held up his index finger: "The only exception!" The companies in the industry form a continuum. Just as it's not always clear where Zeiss ends and ASML begins, the border between ASML and T.S.M.C.

is fuzzy. The chipmaker's relationships with customers are similarly interlaced and sensitive. Two weeks earlier, Cheng told me, he'd hosted representatives from Apple. He said, "I try not to show too much to a customer. I show them as little as possible, because they always want more control."

For T.S.M.C., a new fab isn't just a giant, particle-free building with a twenty-billion-dollar price tag. It's an additional node in a network that extends across space and time. AP7's employees were camped out in temporary offices and bought snacks from a temporary 7-Eleven. Greater energy appeared to be diverted to a construction project across the street: another science-park high school, where the next generation would be educated to take over where the current one leaves off.

In 2023, President Joe Biden opened his State of the Union address by declaring that an injury had been done to our national pride. "These chips were invented in America," he said. "We lost our edge." An injury had also been done to the price of refrigerators and sedans. The remedy, the President said, was that "we're going to make sure the supply chain for America begins in America." The CHIPS and Science Act appropriated fifty billion dollars for the reindustrialization of the country on the back of the semiconductor. The second-largest award went to T.S.M.C., to build out a satellite complex in Arizona.

This was easier said than done. Unlike its kin in Taiwan, the T.S.M.C. fab in Phoenix took about four years to be completed. Permitting and regulations presented the usual challenges, but what seemed to slow down T.S.M.C. the most was that it couldn't rely on the construction pipeline it had cultivated at home. The first time a new contractor works for T.S.M.C., it's given a relatively simple assignment, like a utilities building. The contractor is evaluated according to a set of standardized criteria, and if it does well it graduates to the next level. Cheng, the veteran deputy director, told me, "We do not make bets, and we do not take risks."

The riskiest part, for a company with universally stringent expectations, was hiring. About half the initial workforce was sent over wholesale from Taiwan. During my visit, nobody wanted to be impolitic about the talents of my countrymen, but it was clear that they were dubious about our skill sets and work ethic.

I met Renee Tseng, a young woman in casually elegant denim, who was part of the first wave of Taiwanese employees sent to set up the Arizona operations. She works in the Yield Enhancement Department: before a wafer is routed from one stage of the process to the next, it is scanned by an electron microscope and assessed by an A.I. model, and any detected anomalies pop up on her screen for evaluation. (The company's usage of A.I. effectively requires that it buy back its chips at a retail markup.) She had no formal training; in college, she studied foreign languages. But she's been doing the job for sixteen years, which means that she can immediately distinguish between the curve-shaped scratch made by one tool and the star-shaped scratch made by another. Some people, she said with a shrug, have no aptitude for it: "They all look the same." Others learn the trade, she continued, but it wasn't exactly something that could be taught explicitly. As with many aspects of semiconductor manufacturing, there was a tacit knowledge that could be absorbed only by becoming part of a community of experienced practitioners. Hopefully, in time, her American counterparts would get their reps in.

Another tricky thing about the Arizona site was that it had been severed from its links to local merchants. On one of my fab tours, I'd seen employees push handcarts outfitted with special shock absorbers to protect components from vibrations. These, as far as Jack knew, were made by only one company, in Taiwan, and had to be shipped to the States. There was no clean-room laundry in Phoenix that met T.S.M.C.'s standards, so all the bunny suits were schlepped back and forth across the desert.

The catch, in the American attempt to make chips at home rather than overseas, is that chips aren't really made anywhere at all. A chip can cross international borders

more than seventy times before it's shipped to a customer. T.S.M.C.'s part begins with silicon wafers, but these have their own far-flung prehistory. Leading-edge semiconductor-grade polysilicon must be purified to "eleven nines" (99.999999999 per cent), which only a handful of specialized companies—based in Michigan, Germany, and Japan—are poised to do. These silicon chunks are melted in bespoke crucibles made of the world's purest quartz, virtually all of which is mined from a solitary deposit in the mountains of North Carolina. This process, which generates perfectly crystalline silicon ingots to be sliced like fragile metallic salamis into thin wafers, is entrusted largely to one firm in Taiwan and two in Japan. Most of the steps along this chain are controlled by monopolies or duopolies. This isn't because of market manipulation. It's because each endeavor is so esoteric, so formidable, and so crucial that it makes sense for all the available expertise to be concentrated in one or two nexuses.

Just as important as the firms themselves are the associations between them. Toward the end of the Second World War, as the Red Army approached, the Allied powers moved Zeiss—even then, its lenses were strategically important—from its historic home in Germany's east to the other side of the country. Considered in isolation, the idea of securing some of T.S.M.C.'s operations out of harm's way, in Phoenix, makes similar sense. Broadly speaking, however, it's incoherent to consider any given chip-industry node in isolation. The semiconductor supply chain is a planetary superorganism with its own ecology. It's not like a plant that can be casually repotted.

Chips are products of an organizational culture—both within companies and between them—that encourages people to look beyond their own personal horizons. Jayson Stewart, a San Diego physicist who works on ASML's light source, told me

that the firm's "time frames are very forward-looking. I'm somewhat young, and we're talking about 2050, things to come that are likely past my retirement." There are quite a lot of ASML employees, he continued, who measure their monastic tenure in decades. Some of them have sat in the same clean room in the same bunny suit working on the same little nozzle for fifteen years. He recalled, "When one guy retired, he said, 'Take care of my ASML'—not 'ASML' but '*my* ASML.' " To Schafgans, the engineering executive, even this felt too limited: "I'm grateful to be part of the whole chain." This might sound like brochure patter, but if so it's not the self-promotion of an individual company but the gestalt of the entire industry.




*"It's all vegan, except for some of it."*

Cartoon by Bruce Eric Kaplan

Open cartoon gallery

When I spoke with Mike Schmidt, who ran the main Biden Administration team responsible for the implementation of the CHIPS Act, I expressed some skepticism

about the idea that a domestic industry could be conjured ex nihilo in the Arizona desert. I mentioned the lack of a clean-room laundry, and he said, "I love that about the bunny suits! It just kind of goes to show how much the ecosystem matters, right? Part of it is that in order to do it you have to do it, and then it becomes self-reinforcing. Five years from now, someone will be making a killing with a clean-room-laundry business in Phoenix." A plant or two in Arizona wasn't the U.S.'s ultimate goal but, rather, an initial seed. In time, he hoped, the desert would bloom. This was, he said, "quite self-consciously a generational program."

The best Chinese A.I. models lag behind our own by only three to six months. The near-parity in technical capabilities means little, however, as long as we retain our prodigiously advantaged access to chips. China lacks the virtuosity of T.S.M.C.'s process engineering, and it cannot yet match the lithographic capacity of even ASML's older technologies, let alone its EUV machine. The country will catch up eventually. The rationale for the current export-controls regime is that we might in the meantime build an insurmountable lead.

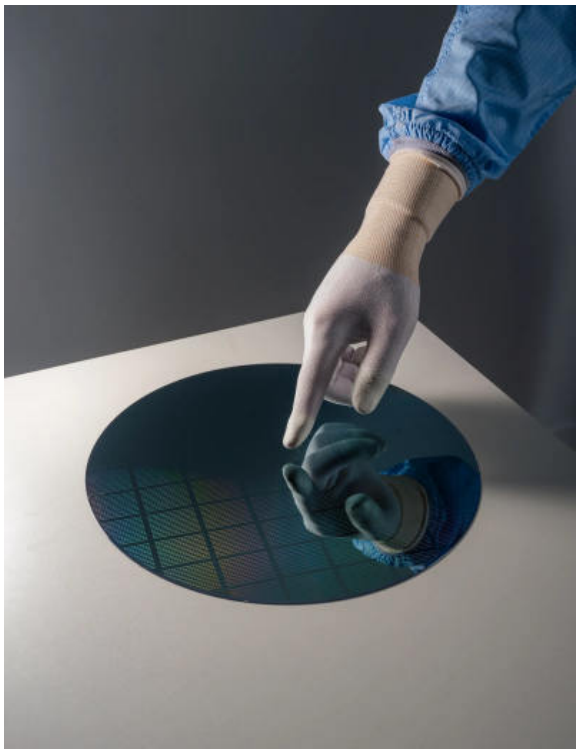
The U.S. government has prohibited the sale of ASML's flagship product in China, but it continues to export an earlier generation of its machines. The Trump Administration has expressed a strong preference for ASML to stop doing even this much. If the company chooses to obey, the Chinese might retaliate with the denial of crucial rare-earth metals.

Europe does not like being caught in the middle. In one recent scenario, European researchers ran a thought experiment in which the Pentagon ultimately identifies the loss of control over ASML's machines as "a threat comparable to the proliferation of nuclear weapons." Washington might then try to impose a joint-governance structure

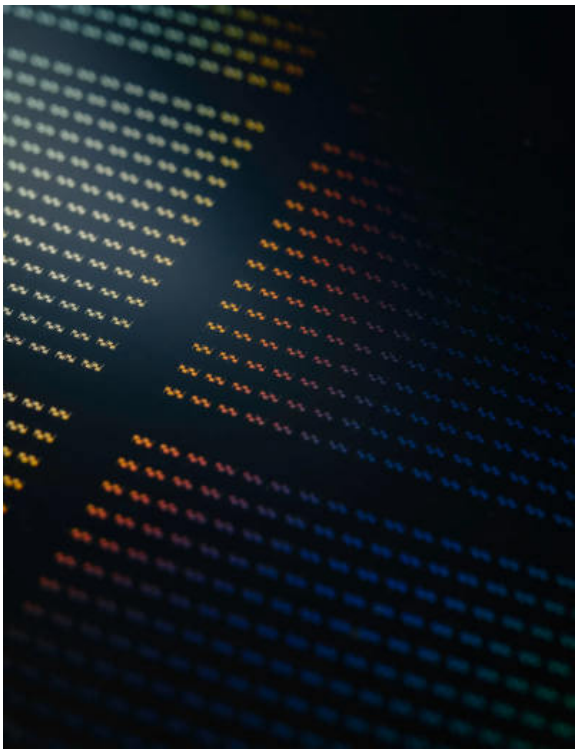
that would effectively nationalize a foreign company. Over dinner, I asked Benschop, the ASML vice-president, about the company's status as a "choke point." He said, "When I learned that word, I thought, I don't want to be choked."

There is no shortage of confident soothsaying about A.I. It will cure cancer; it will replace all jobs; it will transform the battlefield; it will guarantee our political dominance; it will usher in a golden age of superabundance; it will bring about our extinction. What such predictions have in common is an assumption that A.I. developers will enjoy ever-greater access to ever-better chips. That access might be undermined by China. This is a question of geopolitics, in the normal sense of the term. It might also be undermined by the exhaustion of Moore's Law. This, too, might be construed as a question of geopolitics. Moore's Law has survived only on the basis of unprecedented international confederation.

When elected officials or pundits talk about chips, they sound like they could be talking about oil or any other strategic asset that can be extracted. When people who know a lot about chips talk about them, they sound like they are talking about sacramental totems. The first point they want to make is always the same; as Dylan Patel, who runs the research-and-consulting firm SemiAnalysis, told me, "The main thing people don't understand is that semiconductors are the most complicated things people do." The second point they make, before there is time to articulate an objection, is that everyone's intuitive response is incorrect. Patel continued, "Rockets are simple." Brinton Johns, an investor at N.Z.S. Capital, said, of making chips, "It's just so much harder. Rockets are an amazing accomplishment, yes. We don't have a rating scale for complexity, but rockets are maybe a five or a six." His investing partner, Jon Bathgate, finished the thought: "Semiconductors are an eleven."



A silicon wafer printed with ASML's EUV machine. Photograph by Alastair Philip Wiper for The New Yorker



A closeup of a silicon wafer. Photograph by Alastair Philip Wiper for The New Yorker

And yet we take it for granted that we can build them. We shouldn't. The chief determinant of that unlikely success has been not individual genius in the lab, or individual prowess in the boardroom, but collective planning and coördination at a breathtaking scale. Globalization has become tainted in our eyes, as if it necessarily constitutes a zero-sum enterprise. And its mercenary form—which is also its most common form—has indeed inflicted harm. But the global semiconductor supply chain reflects an alternative pathway. It has united its participants in the pursuit of something profoundly difficult that could not have been achieved in the absence of intercontinental solidarity.