Chip War: The Fight for the World’s Most Critical Technology — 20-Minute Read Summary

By Chris Miller (2022)



The Book in 60 Seconds

Every phone, car, hospital machine, and AI system on earth runs on semiconductors — tiny chips whose most advanced versions are made almost entirely in Taiwan, by one company, TSMC, using machines that only one Dutch company, ASML, can build. Chris Miller's *Chip War* tells how that astonishing concentration came to be, from the feuding pioneers of Silicon Valley through Japan's 1980s challenge to the rise of Taiwan's foundry model, and explains why the US–China rivalry over chips is now one of the defining geopolitical contests of our time. It is a business history, a technology story, and a warning: the foundation of the modern economy rests on a single island, and a crisis there would be felt everywhere, immediately.

The Story

The story begins in the late 1940s at Bell Laboratories in New Jersey, where William Shockley helped invent the transistor — the solid-state replacement for the fragile vacuum tube. The transistor could switch and amplify electrical signals while being small, reliable, and cheap to make in volume. It was one of the great inventions of the twentieth century, and Shockley, brilliant and almost universally described as a terrible manager, could not hold onto the people who would make it matter.

In 1956 Shockley moved to Mountain View, California, and founded Shockley Semiconductor Laboratory. Within a year, eight of his best young engineers — the "traitorous eight" — walked out and founded Fairchild Semiconductor. Two of them, Robert Noyce and Gordon Moore, would shape the next half-century of technology. At Fairchild, Noyce developed a practical version of the integrated circuit in 1959, putting many transistors on a single piece of silicon, while Jack Kilby at Texas Instruments arrived at a similar idea at almost the same time. Jean Hoerni's planar process at Fairchild made it possible to build transistors flat on a wafer's surface and connect them reliably — the manufacturing trick that turned the integrated circuit from a laboratory demonstration into a product.

In 1965, Gordon Moore observed that the number of transistors on a chip seemed to double roughly every year — later revised to about every two years — at roughly constant cost. Moore's Law was not a law of nature; it was an observation that became a self-fulfilling prophecy, as the entire industry organized itself around keeping it true. Over fifty years of doubling, chips went from a handful of transistors to tens of billions, each measured in nanometers, billionths of a meter.

Fairchild itself faded as a leader, but it functioned like a university for the industry: its engineers left and founded company after company across the Santa Clara Valley. In 1968, Noyce and Moore left to found Intel with Andy Grove. Intel's first products were memory chips, and in 1971 it introduced the first commercial microprocessor, the Intel 4004, putting a computer's central processing unit on a single chip. The personal computer revolution followed, and the combination of Intel chips and Microsoft software dominated computing for decades.

But leadership in chips has never been permanent, and the first great challenger was Japan. In the 1970s and 1980s, Japanese firms such as NEC, Hitachi, and Toshiba produced memory chips — DRAM — of superb quality and at enormous scale, backed by coordinated government support. American companies lost market share and money; Intel itself exited the memory business in the 1980s to focus on microprocessors, a decision Grove later described as a strategic inflection point. Washington responded with trade pressure, including the 1986 US–Japan semiconductor agreement, and with Pentagon-funded research, because the military understood early that chips were a national security matter, not just a commercial one.

The deeper structural change of that era was the split between designing chips and making them. As factories — fabs — grew more expensive, eventually costing billions of dollars each, a new model emerged: fabless companies that designed chips and paid someone else to manufacture them, and foundries that manufactured chips for many customers without designing their own.

The person who saw this most clearly was Morris Chang. Born in China, educated in the United States, and a senior executive at Texas Instruments, Chang understood manufacturing from the inside and understood trust from the outside: a foundry would only work if it promised never to compete with its customers. In 1987, with backing from the Taiwanese government and the Dutch company Philips, he founded Taiwan Semiconductor Manufacturing Company (TSMC) in Hsinchu, Taiwan. TSMC would make chips for anyone and design chips for no one.

Taiwan was an unlikely setting — a small island with limited natural resources and a contested political status — but it had a government willing to invest in technology, the Hsinchu science park, and a growing pool of well-educated engineers, many US-trained. TSMC became relentlessly good at the hard part of chipmaking: yield, the share of working chips on each silicon wafer, and process technology. More customers meant more revenue, which funded better factories, which attracted more customers — a flywheel that is very hard to stop.

Meanwhile, the American companies that went fabless could focus entirely on design. Qualcomm designed mobile chips, Nvidia designed graphics chips that later proved ideal for artificial intelligence, AMD eventually spun off its factories, and Apple designed its own chips for the iPhone and its computers — all of them manufactured by TSMC. Intel kept its own fabs, and for decades that was an advantage. Then in the 2010s Intel stumbled: it missed mobile, delayed new manufacturing processes, and fell behind TSMC and Samsung of South Korea. Samsung is the other great manufacturer, but it also designs its own chips and phones, so it has never been as trusted a foundry partner as TSMC.

The second great choke point sits in the Netherlands. ASML is the only company in the world that makes extreme ultraviolet (EUV) lithography machines — the tools that print the finest features onto silicon. Lithography works by shining light through a stencil-like mask onto a light-sensitive coating on a wafer, then etching away what is not needed, repeated dozens of times, layer by layer. The smaller the features, the shorter the wavelength of light required, and EUV light is so short it is nearly an X-ray. An EUV machine costs well over $100 million, contains more than 100,000 parts, and uses mirrors polished to atomic smoothness and a light source made by blasting droplets of tin with a laser tens of thousands of times a second. ASML's own supply chain spans the world — mirrors from Germany, light-source technology from an American company it acquired, components from across Europe, the US, and Japan.

The result is a system no country can replicate alone. The United States designs many of the best chips and makes the best chip-design software. Taiwan manufactures them. The Netherlands makes the key machines. Japan supplies critical chemicals and materials. South Korea makes memory. Germany makes precision optics. Remove any one piece and the whole system stalls.

That interdependence held while the great powers tolerated each other. It began to break down as the United States and China came to see each other as rivals — and that is where the “war” in the title begins.

Key Ideas

Moore’s Law is a coordination device, not a law of nature

Moore’s observation — transistor counts doubling roughly every two years — endured because the industry treated it as a target. Roadmaps, investment plans, and research agendas were all built around keeping the doubling going. The lesson is broader than chips: when an entire industry commits to a shared expectation, the expectation can sustain itself for decades. But it also means the system is fragile. Each doubling gets harder and more expensive, and the day the doubling stops, the economics of the whole industry change.

The fabless–foundry split changed who holds power

For decades, a chip company designed and manufactured its own products. The split separated the two, and power migrated to whoever was best at each half. Design became faster and more creative, because fabless companies could focus entirely on architecture and software. Manufacturing became more concentrated, because only a few companies could afford fabs costing many billions of dollars — and only one, TSMC, combined scale, yield, and neutrality toward its customers. The key insight: the most valuable position in the chain is not always the most visible one. Apple and Nvidia are famous; TSMC is, by comparison, quiet — and indispensable.

Trust is a manufacturing technology

TSMC's founding promise — never design its own chips, never compete with its customers — sounds like a legal nicety. Miller shows it was a core technology of the business model. Without it, Qualcomm or Apple would never have handed their crown-jewel designs to an outside factory. In an industry where a leaked design can destroy a company, trust is infrastructure, and it took decades to build. The same logic explains Samsung's limits as a foundry: because Samsung designs phones and chips of its own, customers always wonder whether their designs are truly safe.

Concentration is efficient — until it is fragile

The global chip supply chain is a marvel of efficiency. Each region does what it does best, and the system produces chips of staggering complexity at prices that make a supercomputer affordable in a pocket. But efficiency and resilience are different things. The chain has single points of failure at every level: one company for the most advanced manufacturing, one company for the machines that enable it, a handful of suppliers for key chemicals. The book's central argument is that we optimized for cost and performance for forty years and underpriced the risk of concentration — and that the bill for that risk is now coming due.

Chips are the new oil — and the new arms race

Miller draws the comparison carefully: like oil in the twentieth century, chips are an input to everything, and control over them confers economic and military power. Advanced chips power artificial intelligence, and AI has direct military applications — surveillance, cyber capabilities, weapons systems, logistics. That is why export controls have become the weapon of choice: denying an adversary the most advanced chips, and the equipment to make them, is a way of shaping the military balance without firing a shot. It is also why the analogy has limits. Oil is consumed; chips are durable. Oil comes out of the ground; chips come out of factories that can, in principle, be built anywhere — at enormous cost and over many years.

Export controls work through choke points

The modern US controls on China do not try to block every chip. They target the narrowest parts of the supply chain: EUV machines, advanced chip-design software, and the most advanced chips themselves. Because the chain is so concentrated, a few well-placed restrictions can slow an entire national program. The book is clear-eyed about the costs, though. American chip companies lose real revenue in China and have lobbied for narrower controls. And controls create a powerful incentive for China to build its own capabilities — the question is not whether China will try, but how long it will take and how close to the frontier it can get.

The CHIPS Act is a bet on reshoring — with honest limits

Passed in 2022, the US CHIPS Act provides tens of billions of dollars in subsidies and incentives to build fabs in America. TSMC is building in Arizona, Intel is building in Arizona and Ohio, Samsung in Texas, and Micron is planning a large complex in New York. Miller welcomes the investment but is frank about its limits: a fab takes years to build and longer to run at high yield; the deepest manufacturing expertise and supplier ecosystem remain in Taiwan; US production costs more; and America faces a shortage of trained chip engineers and technicians. Reshoring reduces risk. It does not eliminate it.

The Taiwan Strait is the world’s most dangerous economic fault line

Taiwan is claimed by China, which has never renounced the use of force. It is also home to TSMC and therefore to the world’s most advanced chipmaking. If a blockade or conflict disrupted those fabs, the damage would not be limited to the newest smartphones. Cars, servers, medical devices, military equipment, and the data centers that run AI would all be hit. Other manufacturers could not pick up the slack quickly, because they lack the capacity and the frontier technology. Rebuilding elsewhere would take years and hundreds of billions of dollars, with economic damage measured in the trillions. Miller does not predict war, but he insists the vulnerability is real and cannot be wished away by policy. Taiwan’s dominance is sometimes called its “silicon shield” — the idea that the world has too much at stake to let Taiwan’s fabs stop — but a shield is a deterrent, not a guarantee.

Founders matter, and so do the people who leave

A recurring pattern runs through the book: the people who invent something are often not the people who profit from it. Fairchild trained a generation and then watched its engineers leave to start rivals. Shockley's cruelty scattered his team, and that scattering seeded Silicon Valley. Intel's rise, Japan's challenge, TSMC's founding — each turns on individuals making bets with incomplete information. The lesson for readers of business history is that culture and talent flows matter as much as strategy documents. Silicon Valley exists where it does largely because knowledge walked out the door with people, over and over again.

The AI boom raised the stakes after the book was written

Miller published in 2022, just before large language models made advanced chips the most sought-after industrial product on earth. The AI boom runs on exactly the chips this book is about: Nvidia's graphics processors, manufactured by TSMC on its most advanced processes, and the high-bandwidth memory that sits alongside them. Training a frontier AI model now requires tens of thousands of these chips, and the data centers that house them consume electricity on the scale of small cities. The result is that demand for TSMC's leading-edge capacity has exploded, lead times have stretched, and the strategic value of every advanced fab has risen further. The book's argument — that control over the frontier of chipmaking is control over the future — has been validated by events, but also complicated by them: the AI boom has made the world even more dependent on the very concentration Miller warns about, and it has made the race to build alternative capacity even more urgent. Any reader coming to the book today should hold that post-2022 context in mind: the history is unchanged, but the stakes are higher than they were when Miller finished writing.

The Fight Today

The rivalry the book describes has not cooled since publication — it has intensified, and it now has a clear shape. On one side, the United States has tightened export controls in successive rounds, extending them from the most advanced chips to the equipment and software used to make them, and to the subsidiaries and partners that might serve as conduits. The Netherlands and Japan, home to ASML and to critical materials and equipment suppliers, have been brought into the controls regime, because American restrictions alone would leak. The logic is consistent with Miller's account: choke points only work if everyone who controls one cooperates.

On the other side, China has responded with a whole-of-government push for self-sufficiency. Beijing has poured money into domestic champions, most notably SMIC, and into efforts to build homegrown lithography equipment. Progress has been real but uneven: Chinese firms have made genuine advances in older-generation chips and in some advanced packaging techniques, and SMIC has produced surprisingly capable chips despite the controls, but the frontier — the most advanced processes, at the smallest scales — remains difficult to reach without EUV machines. China has also used its own leverage, restricting exports of gallium and germanium, two materials used in chipmaking, and investing heavily in the open-source RISC-V chip architecture as a way around American-controlled designs.

The third party in the fight is Taiwan itself, and its position has grown more delicate. TSMC's decision to build fabs in Arizona, Japan, and Germany — encouraged by subsidies and by customer demand for geographic diversification — is a hedge against concentration risk, but it also raises a hard question in Taipei: if the most advanced manufacturing spreads abroad, does Taiwan lose some of its silicon shield? Taiwanese leaders insist the most advanced processes will stay at home, and so far that has held. Meanwhile Taiwan continues to invest in its own defense and in its relationship with Washington, while managing an economy deeply intertwined with the mainland.

For the reader, the practical takeaway is that the news will keep coming in waves — a new round of controls, a new fab announcement, a breakthrough claim from a Chinese firm, a delay in Arizona — and the book gives you the framework to sort signal from noise. Ask three questions about any development: does it change who controls a choke point, does it change the geography of the most advanced manufacturing, and does it change the timeline for self-sufficiency on either side? If the answer to all three is no, it is probably noise. If the answer to any one is yes, it matters.

How to Approach It

Do not be intimidated by the technical passages. Miller is a historian, not an engineer, and he explains the technology in plain language. You need only a few concepts to follow the whole book: a transistor is a tiny switch; an integrated circuit packs many transistors onto one piece of silicon; lithography is printing with light; and smaller features require shorter wavelengths and harder machines. Everything else is story.

Read it in two sittings if you can. The first half, up to the rise of TSMC, reads like a business history full of founders, rivalries, and bankruptcies. The second half reads like a geopolitical thriller full of export controls, sanctions, and summit meetings. A pause between the two lets each half land.

Keep a mental map as you go. The cast is large, and companies come and go, but the geography is simple: design in the United States, manufacturing in Taiwan and South Korea, machines in the Netherlands, materials in Japan, assembly in China. Every policy argument in the book is really an argument about one piece of that map.

Finally, read it with healthy skepticism about predictions. Miller is careful not to forecast the future, and you should be careful not to read forecasts into the book that are not there. It explains why Taiwan matters and why a crisis would be catastrophic. It does not tell you whether a crisis will happen — that depends on decisions by leaders in Washington, Beijing, and Taipei that have not yet been made.

Memorable Lines

Miller's prose is clear and journalistic rather than aphoristic, so the memorable lines are less often single sentences than crystallized ideas. A few passages capture the book's argument in miniature:

- On the stakes: the book's central image is that the world's most critical technology is made in one place, by one company, on one island — and that the whole global economy depends on that fact holding.
- On Moore's Law: Miller presents it not as a law of physics but as a shared industry commitment — an observation that became a target, and a target that organized decades of investment.
- On TSMC's founding bargain: the promise never to compete with its own customers was the decision that made the foundry model possible — trust as a business model, not just a virtue.
- On the Taiwan dilemma: Taiwan's chip dominance is described as a "silicon shield" — protection through indispensability — while Miller is careful to note that a shield is a deterrent, not a guarantee.

Fair Criticisms

It is American-centered. The book is written from the perspective of Washington and Silicon Valley, and American interests are the natural frame. That is understandable, but it means Taiwan, Japan, South Korea, and China sometimes get less space than they deserve. Taiwan in particular is treated mostly as a strategic asset — the place where the chips are made — rather than as a society with its own debates, fears, and domestic politics. Readers who finish the book wanting to understand Taiwan itself will need to read further.

Some players are thinly treated. Samsung is a vast and complex company — memory, foundry work, phones, displays — and it deserves more attention than it gets. Japan's materials and equipment firms, which make the chemicals and precision tools everyone depends on, are also somewhat in the background, even though they are choke points in their own right.

The "war" framing cuts both ways. The stakes are genuinely warlike: chips are essential to military power, and a conflict over Taiwan would be catastrophic. But most of the chip war so far has been fought with export licenses, subsidies, and investment decisions, not weapons. Some readers will feel the title makes a complex economic rivalry sound more military and more inevitable than it is. Others will feel the opposite — that given the military applications and the Taiwan scenario, "war" is exactly the right word. The title works as a provocation; the reader should ask whether the framing clarifies or narrows the story.

It is already dated on the details. Published in 2022, the book predates the full force of the AI boom driven by Nvidia's chips and large language models, the further tightening of export controls, and the real-world growing pains of the Arizona fabs. Its central argument has only become more relevant, but the policy chapters should be supplemented with recent reporting.

It is a supply-side book. Miller focuses on who makes the chips and how, and says less about the demand side — the companies and people who build products on top of them. That is a deliberate choice, and it fills in the missing piece beneath the product books in this series, but readers coming from those books will notice the difference in altitude: this is the factories, not the apps.

Why It’s on the List

This is a book about the foundation underneath everything else in this series. The product-management books we have covered — *Hooked*, *Crossing the Chasm*, *The Mom Test*, *The Lean Startup*, *Inspired*, *Continuous Discovery Habits* — all assume a working device in the user’s hand. *Chip War* explains why that assumption holds, how fragile it is, and what it took to build it. It won the Financial Times Business Book of the Year in 2022 because it does something rare: it makes a story about factories and supply chains read like a thriller, and it connects a piece of technology most of us never see to the largest questions of war, peace, and prosperity.

It also belongs on any serious reading list because it changes how you see the news. Once you understand the map — design in America, manufacturing in Taiwan, machines in the Netherlands — every headline about export controls, fab openings, or tensions in the Taiwan Strait becomes legible. You stop seeing a series of disconnected announcements and start seeing a single system under strain. For a reader who wants to understand the world as it actually works, rather than as it appears on a screen, that is an invaluable shift.

Finally, it is a reminder that technology is not magic and not inevitable. The chip exists because specific people made specific bets, failed, quarreled, left, and started again. The system we depend on is a human achievement, built over seventy years, and it is also a human vulnerability. Holding both of those truths at once is what the book asks of you, and it is worth the effort.

What to Read Next

- **The Innovator's Dilemma, by Clayton Christensen** — the classic account of why great companies fail when disruptive technology arrives. It is the perfect lens for Intel's stumble, and for asking whether TSMC itself could one day be disrupted by a new way of making chips. It is coming up in this series.
- **The Design of Everyday Things, by Don Norman** — the other end of the journey. *Chip War* explains how the chip gets made; Norman explains what happens when it ends up in a product in someone's hands, and whether the person using it can figure out how it works. Also coming up in this series.
- **The Cold War, by Odd Arne Westad, or any good history of US–China relations** — for readers who want the wider geopolitical context behind the export-control fight and the Taiwan question.
- **Recent reporting on semiconductors** — because the story moves fast. The AI boom, the tightening of export controls, and the progress of the Arizona fabs have all developed since the book's publication in 2022. Read Miller for the history, then follow current coverage for the present.