Forwards or Backwards

https://forwardorbackwards.com



TSMC: The Company Every Chip Depends On

Watch on YouTube: https://www.youtube.com/watch?v=y6GJmquleEc

TSMC, Taiwan Semiconductor Manufacturing Company, designs nothing it sells under its own name. It builds chips other companies have designed, so it never competes with a customer. This video explains how that one idea, from 1987, made a company most people have never heard of into the factory much of the chip industry depends on.

Morris Chang spent 1958 to 1983 at Texas Instruments, then led Taiwan's Industrial Technology Research Institute, and founded TSMC in 1987. The Computer History Museum calls it a dedicated silicon foundry: an independent factory, available to anyone. In 2025 TSMC says it served 534 customers and made 12,682 different products.

TrendForce's ranking of 9 September 2026 put TSMC at 72.5% of the revenue of the world's ten largest foundries in 2Q26, about USD 40.2 billion, against 5.9% for Samsung and 5.4% for SMIC.

TSMC's own results for 2Q26 show revenue of USD 40.20 billion and a net profit margin of 55.6%. August 2026 revenue was about NT$514.8 billion, up 53.3% on a year earlier. Its 2025 annual filing says high-performance computing was 58% of revenue and smartphones 29%, and that its largest customer paid 19% and its ten largest 78%. The filing names no customers, and neither does this video. NVIDIA names TSMC on its own product pages. Apple's A20 Pro announcement does not.

The newest processes earn most: 3nm and 5nm brought in 30% and 33% of wafer revenue in 2Q26, and 2nm, in volume production since 4Q25, was already 3%. Capital spending was about NT$1.27 trillion in 2025, with USD 52-56 billion planned for 2026.

Most of the most advanced manufacturing is in Taiwan. On 16 July 2026 TSMC announced a further USD 100 billion for Arizona, bringing its planned total there to USD 265 billion: ten fabs, two advanced packaging plants and an R&D centre. The first Arizona fab has made 4nm chips in volume since late 2024; TSMC says its yields are comparable to Taiwan.

The fourth video in the chips series, after CPU nanometres, NVIDIA 4nm and Apple 2nm. Sources are named and dated on screen. Figures that could not be traced to a primary source are left out.

Educational documentary. Not financial or investment advice.

Chapters

1. The name that is not on the box
2. A factory with nothing of its own
3. Why a factory for hire works
4. Seventy-two and a half per cent
5. The money
6. Who pays
7. The newest recipes earn the most
8. The cost of staying ahead
9. Almost all of it in one place
10. Arizona
11. Why nobody simply copies it

1. The name that is not on the box



The last three videos in this series were about chips: how they are made, and what the numbers in their names really mean. Every one of them kept running into the same company, and it is a company most people have never heard of.

Its name is Taiwan Semiconductor Manufacturing Company, usually shortened to TSMC.

Look at how two of the world's best-known chip designers describe their products. NVIDIA's own page says its Blackwell chip is "manufactured using a custom-built TSMC" process. Apple's announcement of its A20 Pro chip says only that it is built on two-nanometre technology, and does not name who built it.

That difference is not an accident, and it is not a secret. It is the business model.

TSMC designs nothing that it sells under its own name. It makes chips that other companies have designed, and puts their names on them, not its own. So it never competes with a single one of its customers.

That one decision, made almost forty years ago, explains nearly everything else in this video.

So this video covers five things. Where the idea came from, and how much of the market it now holds. Who pays for it, where the factories are, and why nobody has simply copied it. The machines inside the factories were covered in the first video of this series, so we will not go back over them here.

We start with a man who spent twenty-five years at an American chip company.

2. A factory with nothing of its own



Morris Chang was born in nineteen thirty-one, in Ningbo, in China. In nineteen fifty-eight he joined Texas Instruments, and he stayed until nineteen eighty-three. By the end he was running the company's worldwide semiconductor business.

In the mid nineteen eighties, the Taiwanese government recruited him to lead its Industrial Technology Research Institute, a government research body. So a career that began in China, and was built in the United States, moved to Taiwan.

In nineteen eighty-seven he founded TSMC, and the Computer History Museum describes what was new about it. It calls TSMC a "dedicated silicon foundry", "an independent factory available to anyone for producing integrated circuits". Integrated circuits is the formal name for chips, and a foundry is simply a factory that makes things to someone else's design.

That was an unusual idea at the time, and it is worth seeing why.

Until then, a company that wanted to sell chips usually had to design them and build them too. That meant owning a factory, which cost a fortune and sat idle whenever its own products did not sell. Chang's model split the two jobs apart: one company designs the chip, and a separate company, with no products of its own, builds it.

Splitting the job that way sounds like a small change. It turned out to reshape the whole industry.

3. Why a factory for hire works



The reason it works comes back to trust.

If you design chips, you do not want your competitor to see your designs. A factory owned by a rival is a risk, however good it is. A factory that sells nothing of its own has no reason to copy you, so you can hand it everything, and so can your competitors.

Because so many companies can share one set of factories, each of them avoids building its own. And the factory stays busy, because when one customer's orders fall, another's rise.

The scale of that sharing is in TSMC's own figures. In twenty twenty-five, it says, it served five hundred and thirty-four customers. Between them they had it make twelve thousand, six hundred and eighty-two different products. It offered three hundred and five different manufacturing processes to make them with.

A process here means a particular recipe for building a chip, tuned for speed, for low power, or for cost.

A company that designs chips but owns no factory is called fabless, from the word fab, which is the industry's name for a chip factory. The fabless company spends its money on designers and software. The foundry spends its money on buildings and machines, and every fabless customer shares the bill.

That arrangement is now the normal way to make an advanced chip. And one foundry makes far more of them than anyone else.

4. Seventy-two and a half per cent



TrendForce, a market research firm, ranks the world's largest foundries every quarter by the money they earn.

Its ranking for the second quarter of twenty twenty-six was published on the ninth of September. It put TSMC's revenue at about forty point two billion dollars. The next largest foundry, Samsung, earned about three point two six billion. China's SMIC was third, at about three billion.

Put those numbers side by side and the gap is hard to take in.

Measured as a share, TrendForce says TSMC took seventy-two and a half per cent of the revenue of the ten largest foundries. Samsung took five point nine per cent. SMIC took five point four per cent.

That figure is often quoted, so it is worth being exact about what it measures.

It is a share of foundry revenue: money earned making chips to order for other companies. Chips that a company makes in its own factories are not part of it. TrendForce put the ten largest foundries' combined revenue at about fifty-three point five billion dollars for the quarter, a record.

Seventy-two and a half per cent of a record is a great deal of money. What does that look like for TSMC itself?

5. The money



TSMC publishes its results every quarter, and its sales every month.

In the second quarter of twenty twenty-six, TSMC reported revenue of forty point two billion US dollars. Of every dollar it took in, fifty-five point six cents was left as profit after all its costs and taxes. That share is called the net profit margin.

The monthly figures are in New Taiwan dollars, and they show the shape of the year so far.

Month by month, twenty twenty-six has mostly climbed. February was the low point, at about three hundred and eighteen billion New Taiwan dollars. August was the highest month of the year so far, at about five hundred and fifteen billion, up fifty-three per cent on a year earlier.

Those figures are the company's own, and it notes that the monthly numbers have not yet been audited.

TSMC also says what it expects next. For the third quarter of twenty twenty-six, it guided revenue of between forty-four point six and forty-five point eight billion dollars. That is a forecast, and it is shown here as the range the company gave, not as a single number.

So where does all of that money come from?

6. Who pays



TSMC's annual report to American regulators, filed in April twenty twenty-six, breaks its revenue down in several ways.

By the kind of product, fifty-eight per cent of its twenty twenty-five revenue came from high-performance computing, which mostly means the chips in data centres. Smartphones were twenty-nine per cent. By region, customers based in North America brought in seventy-five per cent.

It also shows how much depends on a few customers.

The largest single customer paid nineteen per cent of TSMC's revenue in twenty twenty-five. The second largest paid seventeen per cent. Together, the ten largest customers paid seventy-eight per cent of everything.

The report does not say who those customers are, and this video will not guess. What can be said is what the customers themselves say. NVIDIA names TSMC on its own product pages. Apple's announcement of the A20 Pro does not name its manufacturer at all.

A company with that many dollars flowing from so few names has an obvious question to answer: what exactly are they paying for?

7. The newest recipes earn the most



TSMC's results split its wafer revenue by process. A wafer is the thin disc of silicon that chips are made on, as the first video in this series showed.

Look at the second quarter of twenty twenty-six. The three-nanometre process brought in thirty per cent of TSMC's wafer revenue. Five-nanometre brought in thirty-three per cent. Seven-nanometre added another eleven. The newest process, two-nanometre, was already three per cent.

Those names are not measurements. As the first video in this series explained, a nanometre figure in a process name is a product name, not the size of anything on the chip. What matters here is how quickly the newest processes take over the revenue.

Taken together, TSMC's processes at seven nanometres and below earned seventy-seven per cent of its wafer revenue in that quarter.

The two-nanometre process is very new. TSMC's own page says it started volume production in the fourth quarter of twenty twenty-five, as planned. Within about six months, it already accounted for three per cent of wafer revenue.

Staying at the front like that is the most expensive habit in manufacturing.

8. The cost of staying ahead



TSMC's capital spending, meaning the money it puts into new buildings and machines, was about one point two seven trillion New Taiwan dollars in twenty twenty-five. For twenty twenty-six, it planned between fifty-two and fifty-six billion US dollars. Its factories could already produce more than seventeen million wafers a year, counted in the industry's standard twelve-inch size.

The reason the spending never stops is that every new process needs new equipment, and much of the old equipment cannot make the new chips.

A new process needs new factories, fitted with some of the most complex machines ever built. Each one has to be installed, tuned and learned before it produces anything useful. That learning is slow, and it cannot be bought.

It works as a loop. Customers pay most for the newest processes. That money pays for the next round of factories. The next round of factories wins the next round of customers, and the loop turns again.

The loop explains the money. It does not explain the map.

9. Almost all of it in one place



TSMC lists its factories on its own website, and the list is short.

In Taiwan it runs six of its largest factories, which it calls GIGAFABs, plus several older, smaller ones. Outside Taiwan it has one site in Arizona and one in Washington State in the United States, one in Japan, and one in China. A specialty factory in Dresden, in Germany, is being built.

A GIGAFAB is TSMC's word for its biggest factories, each one built to process the standard twelve-inch wafers.

So the most advanced of TSMC's manufacturing sits on one island. That is simply what the company's own list shows. It is also why the expansion in the next chapter matters so much.

The overseas sites are not copies of Taiwan. The Washington State factory works on older, eight-inch wafers, and the German factory is described as a specialty plant. Outside Taiwan, the site that TSMC is building up for its newest processes is Arizona.

So what is actually happening in Arizona?

10. Arizona



The first Arizona factory has been in volume production since late twenty twenty-four, making chips on TSMC's four-nanometre process. The second has finished construction, with volume production expected in twenty twenty-seven. Then, on the sixteenth of July twenty twenty-six, TSMC announced a further expansion.

The announcement was relayed the same day by the Arizona Commerce Authority, the state's economic development agency.

TSMC said it would invest another one hundred billion dollars, for four more factories at two nanometres or beyond. That brings its total planned investment in Arizona to two hundred and sixty-five billion dollars. The plan now covers ten factories, two advanced packaging plants and a research centre.

Advanced packaging is the step that joins finished chips together into one working part.

One claim in that announcement is worth singling out, because it is the one that matters most. TSMC says the first Arizona factory is running "with yields comparable to TSMC's facilities in Taiwan". Yield means the share of chips on each wafer that actually work, so a factory with poor yields throws away much of what it makes.

That claim comes from the company itself. It is attributed here as the company's statement, not as an independent measurement.

On a map, the size of the move is plain. Most of TSMC's largest factories are in Taiwan, and Arizona is on the far side of the Pacific Ocean. The first factory there makes four-nanometre chips, and the four new ones are planned for two nanometres or beyond.

11. Why nobody simply copies it



If the model is so good, why has nobody matched it? The figures in this video suggest four reasons.

The first is trust, because TSMC still competes with none of its customers. The second is scale: five hundred and thirty-four customers keep its factories full. The third is money, spent at a pace few companies can match. The fourth is learning, which can only be gained by running the newest processes for years.

None of those four can be bought quickly. Each one depends on the others, which is why the loop from chapter eight is so hard to break into.

The distance to second place makes the point. In the second quarter of twenty twenty-six, TrendForce put Samsung's share of foundry revenue at five point nine per cent. TSMC's was seventy-two and a half.

That is where the story ends up: a company whose name is on almost nothing, making the chips that almost everything else depends on.

Follow any advanced chip back to where it was made. The name on the box is the company that designed it. The factory behind it very often belongs to TSMC. That split, between the designer and the builder, is the idea Morris Chang started with in nineteen eighty-seven.

Whether that factory can be built again somewhere else is the question the company is now spending hundreds of billions of dollars to answer.

Sources and credits

Primary sources

- TSMC, company profile, tsmc.com - established 1987 by Dr. Morris Chang; pure-play foundry; in 2025, 534 customers, 12,682 products, 305 technologies, annual capacity over 17 million 12-inch equivalent wafers; fab locations in Taiwan, the United States, China, Japan and Germany.
- Computer History Museum, Morris Chang, 2007 Fellow profile - born 1931, Ningbo; Texas Instruments 1958-1983, group vice president for worldwide semiconductors; recruited to lead ITRI; founded TSMC in 1987 as a 'dedicated silicon foundry'.
- TSMC, '2Q26 earnings release', pr.tsmc.com, 16 July 2026 - revenue NT$1,270.38 billion / USD 40.20 billion, up 36.0%; net income NT$706.56 billion; gross margin 67.7%; net profit margin 55.6%; wafer revenue 2nm 3%, 3nm 30%, 5nm 33%, 7nm 11%, 7nm and below 77%; 3Q26 guidance USD 44.6-45.8 billion.
- TSMC, monthly revenue 2026, investor.tsmc.com, and 'August 2026 Revenue Report', pr.tsmc.com, 10 September 2026 - January to August 2026 revenue (unaudited); August NT$514.81 billion, up 53.3%; January-August NT$3,386.87 billion, up 39.3%.
- TrendForce, 'Global Foundry Revenue Approaches US$53.49 Billion in 2Q26', 9 September 2026 - TSMC about USD 40.2 billion, 72.5% share of top-10 foundry revenue; Samsung 5.9%; SMIC 5.4%; UMC and GlobalFoundries.
- TSMC, 2025 Annual Report on Form 20-F, filed with the SEC 16 April 2026 - 2025 net revenue NT$3,809,054 million; largest customer 19%, second 17%, ten largest 78%; HPC 58%, smartphone 29%; North America 75%; 2025 capex NT$1,272,411 million; 2026 capex plan USD 52-56 billion.
- Arizona Commerce Authority, 'TSMC Announcement', 16 July 2026 - a further USD 100 billion for four more 2nm-or-beyond fabs; Arizona total USD 265 billion: 10 fabs, 2 advanced packaging facilities and an R&D center; first fab in N4 volume production since late 2024 with yields comparable to Taiwan; second fab complete, volume production expected 2027.
- NVIDIA, Blackwell and Hopper architecture pages (via the Nvidia 4nm fact sheet) - 'manufactured using a custom-built TSMC 4NP process'; 'the TSMC 4N process customized for NVIDIA'. Apple newsroom, 9 September 2026 (via the Apple 2nm fact sheet) - A20 Pro 'built using the latest 2-nanometer process technology'; TSMC not named.

Not regulated financial advice.