Investor Letter

3rd Quarter 2023

"It's the eye of the tiger,
It's the thrill of the fight,
Rising up to the challenge of our rival,
Stalks his prey in the night,
And he's watching us all,
With the eye of the tiger."

Excerpt from the song "Eye of the Tiger" by American band Survivor.

The SFA EAC FIC FIA – BDR LEVEL I Fund closed the 3rd quarter of 2023 with a year-to-date return of 19.65%, outperforming its benchmark Ibovespa, which returned 6.22%.

Período	SFA EAC	IBOVESPA	CDI
3T23	2,54%	-1,29%	3,22%
YTD	19,65%	6,22%	9,93%
Últimos 3 anos	17,52%	20,17%	29,59%
Últimos 5 anos	99,80%	48,24%	42,70%
Desde o início	311,47%	117,85%	147,08%

As 2023 approaches its end, we reflect on a year filled with surprises and challenges. Defying initial expectations, we observed robust economic growth, not only in the United States but also in Brazil, where growth projections were revised significantly upward.

This year was marked by an environment of accelerated growth accompanied by an expressive disinflation process. Such an economic environment allowed the Central Bank of Brazil to begin an interest rate-cutting cycle, which has direct implications for our equity market.

Over the past two years, high interest rates proved a considerable obstacle for the Brazilian stock market, negatively impacting everything from corporate financing costs to household consumption among highly indebted families. The exit of investors from equities added complexity to the scenario, making market navigation more challenging. However, with the recent decline in interest rates, what was once a headwind has now become a tailwind. This optimism is reinforced by the expectation that the Selic rate will reach a single-digit level.

During 2023, we identified and seized numerous opportunities to allocate capital. We increased our stake in companies already in our portfolio or that we have followed for many years. We also became partners in fantastic new companies. We have invested considerable time deepening our understanding of their business models and value-creation potential.

We remain excited about the prospects of the companies in our portfolio, while always attentive to challenges and scenario changes so that we can adapt with agility in an active portfolio management approach.

Brief History of Computing

Before gaining such importance in the modern world, semiconductors went through a long cycle of discoveries and maturation. The history of computing is a vibrant example of human ingenuity.

In the ancient civilizations of the Babylonians and Egyptians, sketches of algorithms and calculations were already found on clay and stone. Later, in the 9th century in Persia, Al-Khwarizmi wrote the first concepts of algorithms. The root of the words “algarism” and “algorithm” comes from *algoritmi*, the Latin form of his name. In the 17th century, in the context of the industrial mechanical revolution, Pascal and Leibniz forged mechanical calculators.

In the 19th century, Charles Babbage developed the concept of an analytical machine (mechanical computer). This machine of gears and steam would be capable of incorporating an arithmetic logic unit, control flow in the form of conditional branches, loops and integrated memory, making it the first design for a general-purpose computer that could be described in modern terms as *“Turing-complete”*. Despite the concept, its implementation did not emerge until 1940. Meanwhile, George Boole’s algebra was being shaped, governing the very logic of digital circuits. The concept of Alan Turing’s ‘universal machine’ also emerged, a fundamental concept in computer science.

In 1940 the ENIAC computer emerged, designed to calculate artillery firing tables for the U.S. Army Ballistic Research Laboratory. However, its first program was a study on the feasibility of a thermonuclear weapon. We also had the “Enigma”, an electromechanical cipher machine with rotors, depicted in the excellent film *“Imitation Game”*, used to encrypt/decrypt war codes. The Enigma machine was used by Germany to communicate.

In 1947 the first PNP germanium transistor was demonstrated, which enabled the idea of mass-producing transistors. Subsequently, silicon would prove to be a more interesting alternative. In the 1970s and 1980s, the silicon microprocessor rose to prominence, a fundamental development that inaugurated an era of accessibility to PCs.

Finally, we arrive at the modern era, the 2000s and 2010s, where the marvels of mobile computing, cloud computing and social media transform every facet of human existence, and the flourishing of artificial intelligence and *big data* analysis enable tremendous value creation.

Enigma machine (used in WWII)

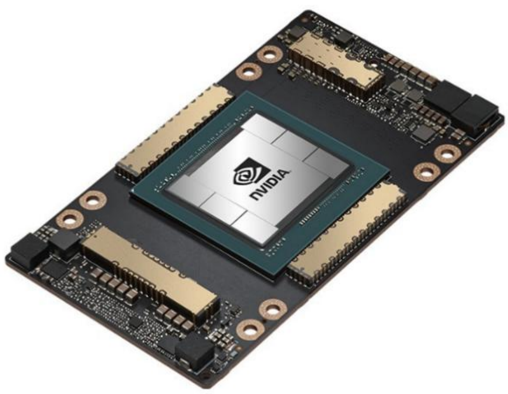

A100 – Nvidia GPU widely used for language model inference

Semiconductors: The Industry That Is the Foundation of the Modern World

The complex semiconductor chip industry is one of the most important for the functioning of the modern economy. If you are using an LED (*light emitting diode*) lamp to light your environment, its entire operation — from the control circuit to the emission of light itself — depends on semiconductors. The vehicle you use for transportation likely depends on semiconductors to perform the most basic functions (combustion system, electrical system, information center, presence sensors, lights, etc.). When we watch a film on our TVs, advanced integrated circuits process all input signals and the operating system, projecting the information onto a *crystal* panel, which to generate an image also contains, again, thousands of transistors. Today, all industries use semiconductor chips at some link in the production chain.

In the semiconductor industry, TSMC is responsible for acting in the so-called*foundry*) of these chips, which is the direct printing of all integrated circuit logic onto the silicon substrate (*Wafer).*

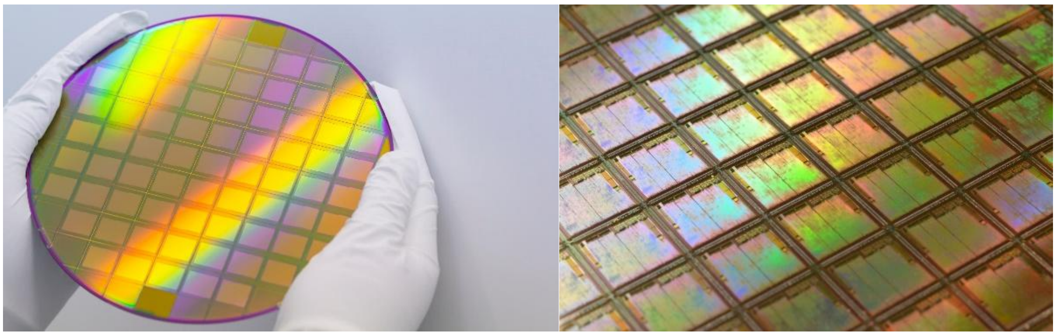

Wafers after microchip production

If you are using a mobile phone to read this text, your device may have more than 20 billion transistors. These are 20 billion components, meticulously designed to operate jointly in a complex logical network, conceived with a very *design* complicated design. It is like imagining a “puzzle” of 20 billion pieces, where in addition to “fitting” with the neighboring piece, each one must send and receive predefined information in the form of electric current. All of this must function perfectly for us to have a modern mobile phone we use for the simplest tasks, such as sending a message or opening a PDF.

To put that number in perspective, if we had 20 billion grains of sand, we would need 20 ten-liter buckets to hold them! How is it possible to carry in our pocket a device with so many components?

Blueprint de um celular moderno – gerado por AI

Representação da quantidade de transistores de um celular moderno em grãos de areia – gerado por AI

So, How Did We Get to 20 Billion Transistors?

All digital electronics are based on a very simple component, the transistor. In a simple analogy, the transistor is to a faucet as the electronic circuit is to a water circuit. The conventional transistor controls the flow of current according to the signal at the base (B in the figure). In this way, flow can be conditionally controlled, just as a faucet allows the passage of water depending on whether its valve is open or closed.

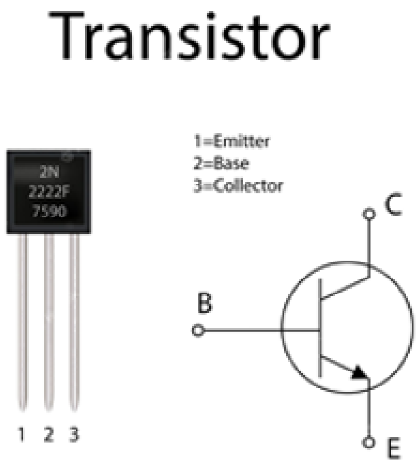

Transistor controls current flow

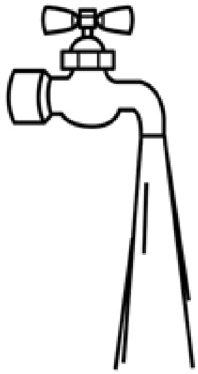

Faucet (or valve) controls water flow

From this basic element we can build everything else. It involves layers of abstraction on top of other layers of abstraction until we reach the CPU. Starting from the transistor we then have logic gates, combinational circuits, arithmetic and logic units, memory, a control unit, and finally the CPU (central processing unit), which revolutionized the world from the 1960s onward.

From this basic component and with great creativity and complex designs, it is possible to perform the most diverse functions that chips perform today — just as the same brick can build many types of structures with distinct styles and purposes.

Moore's Law and "Nanometry"

As new applications for computers emerged and more and more problems migrated from the physical world to the digital world, demand for microprocessors increased. The advent of the internet and the mobile phone era was a new wave of transistor consumption. And how to produce ever more transistors and accommodate them in ever-smaller spaces? Simple: by shrinking the transistors. Moore's Law, theorized in the 1970s, predicts that the number of transistors on a chip will double, at the same cost, every two years. The trend of this observation of the pace of electronic computing evolution results in a lower production cost per transistor on each chip. This reduction in unit cost is necessary to justify new technological generations. Unit cost declined and, as a result, the value per area of a chip increased. A *wafer* (substrate on which chips are manufactured) today is worth thousands of times more than it was 20 years ago.

As transistor size decreased, the distance between transistors on a chip came to be measured in nanometers, and the nanometers indicate the evolution of that processor generation: the smaller, the more modern. Every 1 to 2 years a new generation with smaller “nanometry” is mastered.

Note that the charts below are on a **logarithmic scale**.

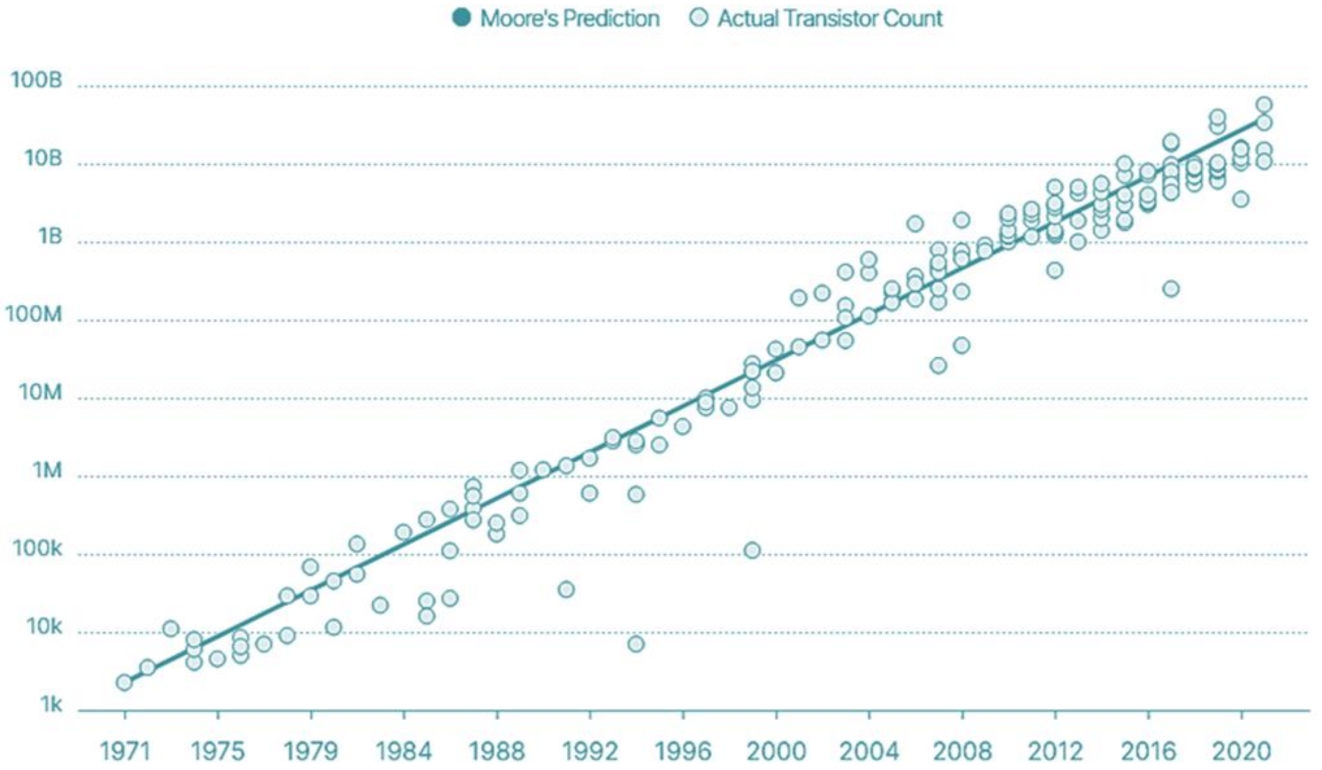

Number of transistors grows exponentially

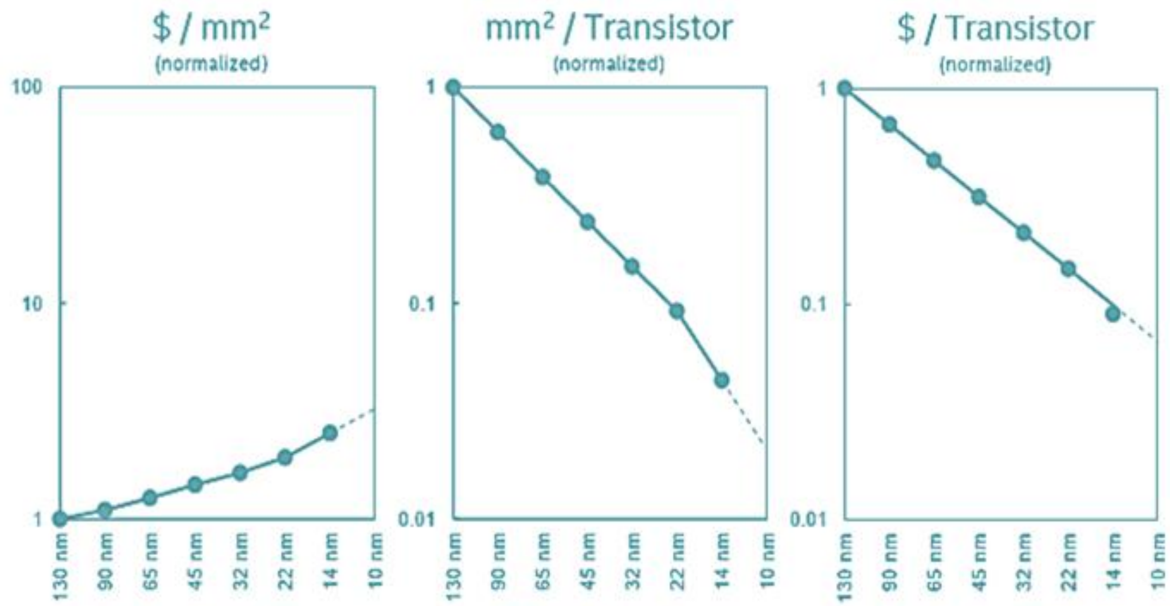

Transistor density per area grows exponentially, and monetary value per area also grows exponentially — TSMC’s revenue follows accordingly.

Nanoscale Manufacturing

To deliver the smallest and most efficient chips in the industry, TSMC is at the heart of what is the most complex manufacturing industry humanity operates today. It is responsible for integrating the entire supply chain. TSMC coordinates with raw material and equipment suppliers, builds factories, and produces chips at high scale and efficiency for the designs requested by its clients. Below, the image illustrates various nanometer sizes. For comparison, a DNA strand is 2.5nm wide/in diameter. The latest iPhone 15 already features, for the first time, 3nm technology.

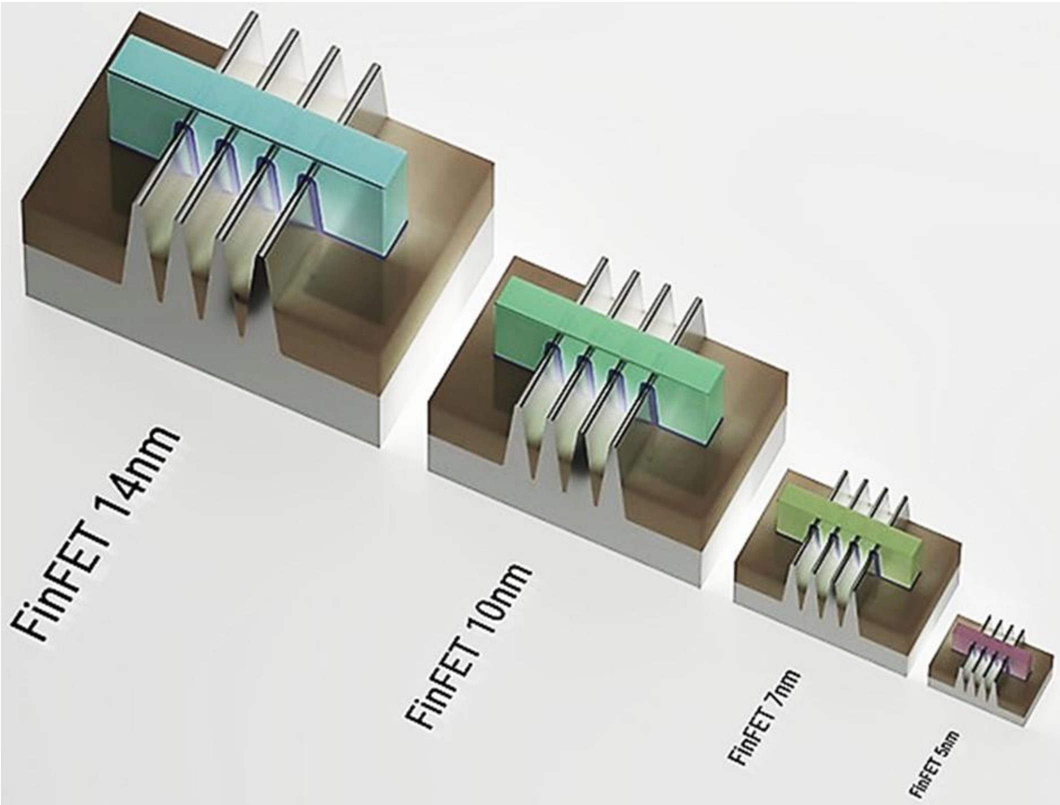

Representation of various sizes: 14nm first manufactured in 2014, 5nm first manufactured in 2020

20 µm – 1968	90 nm – 2003
10 µm – 1971	65 nm – 2005
6 µm – 1974	45 nm – 2007
3 µm – 1977	32 nm – 2009
1.5 µm – 1981	22 nm – 2012
1 µm – 1984	14 nm – 2014
800 nm – 1987	10 nm – 2016
600 nm – 1990	7 nm – 2018
350 nm – 1993	5 nm – 2020
250 nm – 1996	3 nm – 2022
180 nm – 1999	Future
130 nm – 2001	2 nm ~ 2024

Clients Choose TSMC

"Our strategy is very simple. Our strategy is to be the most trusted technology and capacity provider in the semiconductor industry."
Morris Chang, industry veteran and founder of TSMC

The Taiwanese manufacturer holds technological leadership in chip manufacturing, constantly pushing new frontiers in R&D. TSMC's greatest competitive advantage is its irreplicable chip development and manufacturing process. TSMC's unique chip manufacturing process generates efficiency and economies of scale. Since TSMC's founding, efficiency — in the form of reduced production cycle time — has remained a goal pursued to this day. When the industry manufactured a chip in 40 to 50 days, Morris Chang, TSMC's founder, set a goal of manufacturing in a 25-day cycle. By manufacturing a chip in half the time of its competitors, TSMC is able to go through more manufacturing cycles in less time, quickly identify and correct flaws, and optimize its production line. Lithography adjustment processes are iterated more rapidly, and other equipment is also adjusted in less time for each *ramp up* stage of a new chip generation. In this way, Morris ensured that TSMC had less waste. A good portion of the chips on a *wafer* are lost due to quality failures; if a grain of sand enters the *clean room* it can compromise an entire chip. The *yield* de um *Wafer* is defined as the percentage of chips tested and approved for delivery to the customer out of the total manufactured. TSMC has thus always had the best *Yields,* which also contributes positively to their unit cost.

Furthermore, they have very long-standing commercial relationships with their customers — Apple, for example, has produced with TSMC since 2014. In this relationship and with other major customers, we can note a switching cost from which TSMC benefits (although it does not constitute an extremely strong competitive advantage). The switching cost would relate to the deep integration of the entire development cycle of a new product. A customer like Apple or Nvidia, if they ever switched, would need to ensure that all the logistics and complicated *supply chain* functioned well from one year to the next. Given the scale and complexity involved, this is no simple task. TSMC acts as an orchestrator, itself maintaining a large number of suppliers and basic inputs to carry out its activities.

TOP 10 clientes da TSMC

	2015	2020	2022
Apple	15,9%	25,1%	23,9%
AMD	4,3%	7,0%	10,2%
Qualcomm	15,9%	8,0%	8,2%
MediaTek	6,7%	5,9%	8,1%
Nvidia	5,8%	7,1%	7,6%
Broadcom	7,4%	5,9%	6,6%
Intel	4,1%	4,1%	5,0%
Will Semi	2,8%	1,8%	1,9%
NXP	1,7%	1,3%	1,6%
Marvell	3,9%	1,2%	1,6%
Infineon	1,8%	1,3%	1,5%
Hi-Silicon	4,1%	12,5%	0,0%
Soma	74,4%	81,2%	76,2%

Despite a relative concentration of customers, it has very resilient clients such as Apple, AMD, Nvidia, and Intel (even Intel, which also has its own *foundry*).

How Dedicated *Foundries* Emerged

"He who controls the spice controls the universe."
Frank Herbert, "Dune".

The dedicated*foundry*) activity emerged in the 1980s with TSMC, founded in 1987, being one of the first *foundries*, in a plan that was at the time innovative and even counterintuitive: to focus on the part of the chain that only manufactures semiconductors. Until then, the industry operated in a vertically integrated manner, and each player closely guarded its secrets and techniques. Over the years it became increasingly expensive and complex to maintain this industrial foundry park (*foundry*) and each new paradigm represented a significant capital commitment, as well as R&D risks if the technology did not succeed. As an example, in the late 1980s a battle was waged between the technologies "*Photolithography (G-line, I-line)*" and "*X-Ray litography*" and there were always those who bet on one technology over the other, and uncertainty was high. Only through trial and error is the most efficient production method discovered. IBM lost a great deal of money with years of research and investment in "*X-Ray litography*", despite also having achieved success with photolithography in the 1990s.

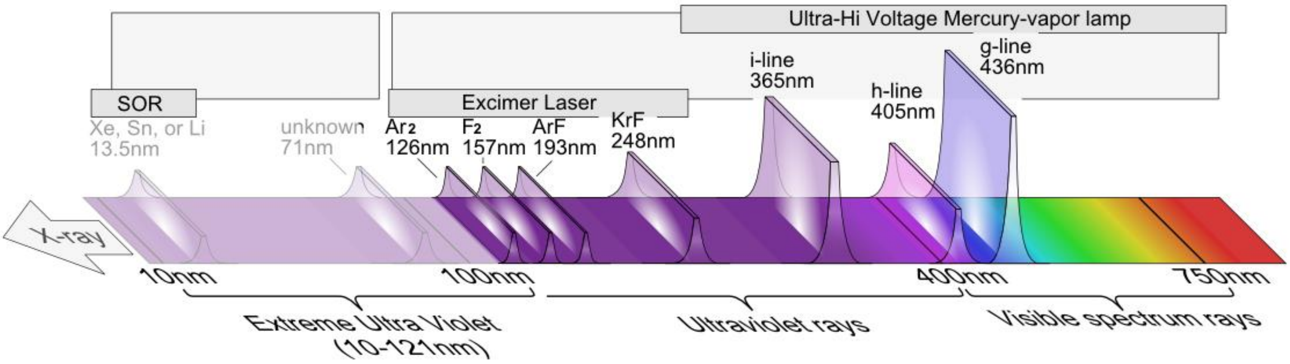

R&D was a long and uncertain process leading to current technologies. The wavelength of light is linked to the possible manufacturing sizes

The message is: it is becoming riskier and more expensive to be at the vanguard, at the frontier of knowledge in terms of microchip manufacturing. Many companies fell by the wayside, unable to raise capital or conducting research but failing to deliver viable technologies. It is an industry that requires a great deal of accumulated knowledge, as well as a strong balance sheet to conduct multiple R&D efforts simultaneously; the fundamental imperative is not to fall behind and always find a way to produce *Wafers* new with *yield* reasonable yield, rapidly ramp up production, and consistently deliver high quality in a customized manner according to each customer's needs.

Therefore, in a natural evolution of the sector, there were many casualties (insolvencies and acquisitions), and those still able to produce chips found themselves under increasing pressure from greater capital and skilled labor requirements to remain in the race. In this environment, during the 1990s TSMC grew consistently, and for its customers outsourcing chip production made more sense as they mitigated their own R&D risks and could deploy more capital into design and product — a trend from which there would be no turning back.

Competition

“When the winds of change blow, some people build walls and others build windmills.”
Chinese proverb.

Over time it became increasingly attractive for most companies to shed their *foundry* (foundry) operations and focus on the commercial aspects of the final product. This allowed many fabless companies to emerge. These are companies focused on the upper part of the chain, such as ARM, which emerged in 1990, or Nvidia, founded in 1993.

In the 1980s as a pure foundry, TSMC had almost no competitors. It was the first company to dedicate itself entirely to *foundry services.* In the 1990s, UMC, also Taiwanese, followed in its footsteps and pivoted to a fully services and manufacturing model. Today UMC, despite being a competitor, holds half of its revenue in nodes larger than 50nm, and its most advanced is 22nm, which is still in the process of being “*ramp up*”. We can see that UMC does not offer a competitive threat to TSMC in the most advanced technologies. In the 1990s it became evident that “*pure foundries”* were here to stay and many companies abandoned the vertically integrated model, realizing it made more sense to outsource this stage of the process. While others bet on specializing in *foundry.*

By the 2000s, most companies had either died or given up on the ambitious undertaking of pursuing high technology in microchip manufacturing. At that point, only 5 companies battled for chip manufacturing proficiency: Intel, Samsung, UMC, GlobalFoundries (AMD), and Chartered Semiconductor Manufacturing. Chartered, however, was sold to GlobalFoundries in 2009, as it was too small and was swallowed by the high investment requirements — its scale fell below the break-even scale and its speed below escape velocity. Despite not having committed any strategic blunder, it is a good example of how scale matters.

In the 2010s, TSMC took absolute leadership of this market, securing important contracts such as manufacturing all Apple iPhones in 2014. Intel, in addition to not having prioritized third-party production, stumbled on essentially technical difficulties and significantly delayed its 10nm technology, falling virtually far behind ever since. Intel always promised but delivered below expectations. Today they understand the need to have the most advanced manufacturing technologies and are investing to close the *gap*. Moreover, Intel is still a significant customer of TSMC itself.

The last to throw in the towel was GlobalFoundries. Despite being an excellent company, its high-technology projects became too expensive and risky. Without a balance sheet as strong as TSMC’s, a first-tier established customer network, and already behind in the race, they ended up abandoning the attempt to continue the war — from a passage in the book *Chip Wars* by Chris Miller, in free translation:

“TSMC, Intel, and Samsung were right to adopt EUV, although they had different strategies about when and how to embrace it. GlobalFoundries was less confident. The company struggled with its 28nm process. To reduce the risk of delays, it decided to license its 14nm process from Samsung rather than develop it internally — a decision that did not suggest confidence in its R&D efforts.

By 2018**, GlobalFoundries had purchased several EUV lithography tools and was installing them at its most advanced plant, Fab 8, when company executives ordered work to stop. The EUV program was being cancelled. GlobalFoundries was giving up on producing new leading-edge nodes. It would not pursue a 7nm process based on EUV lithography, which had already cost US$1.5 billion in development and would require a comparable amount of additional spending to have a production line up and running. TSMC, Intel, and Samsung had financial positions strong enough to bet and wait that they could make EUV work.* ***GlobalFoundries decided that, as a mid-sized foundry, it could never make a 7nm process financially viable. It announced it would stop building ever-smaller transistors, cut R&D spending by one third, and quickly turned a profit after several years of losses. Building leading-edge processors was too expensive for all but the world’s largest

__chip manufacturers.__ Even the deep pockets of the Persian Gulf kings, owners of GlobalFoundries, were not enough. The number of companies capable of manufacturing leading-edge logic chips fell from four to three."

Value Chain

Therefore, in a natural evolution of the sector, there were many casualties (insolvencies and acquisitions), and those still able to produce chips found themselves under ever-greater pressure from increasing capital and skilled labor requirements to remain in the race. In this environment, TSMC flourished and for its customers, outsourcing production became permanent, as they mitigated their own R&D risks and could deploy more capital into design and product — again, a path of no return.

Regarding the supply chain, we draw attention to the diagram below, which shows how American firms have turned toward design and intellectual property. While equipment suppliers are also an important part of the complex semiconductor value chain, TSMC is focused on on-demand manufacturing — a services company since its inception — performing mainly chip fabrication on *wafers* (*front end*) and assembly and test services (*back end*).

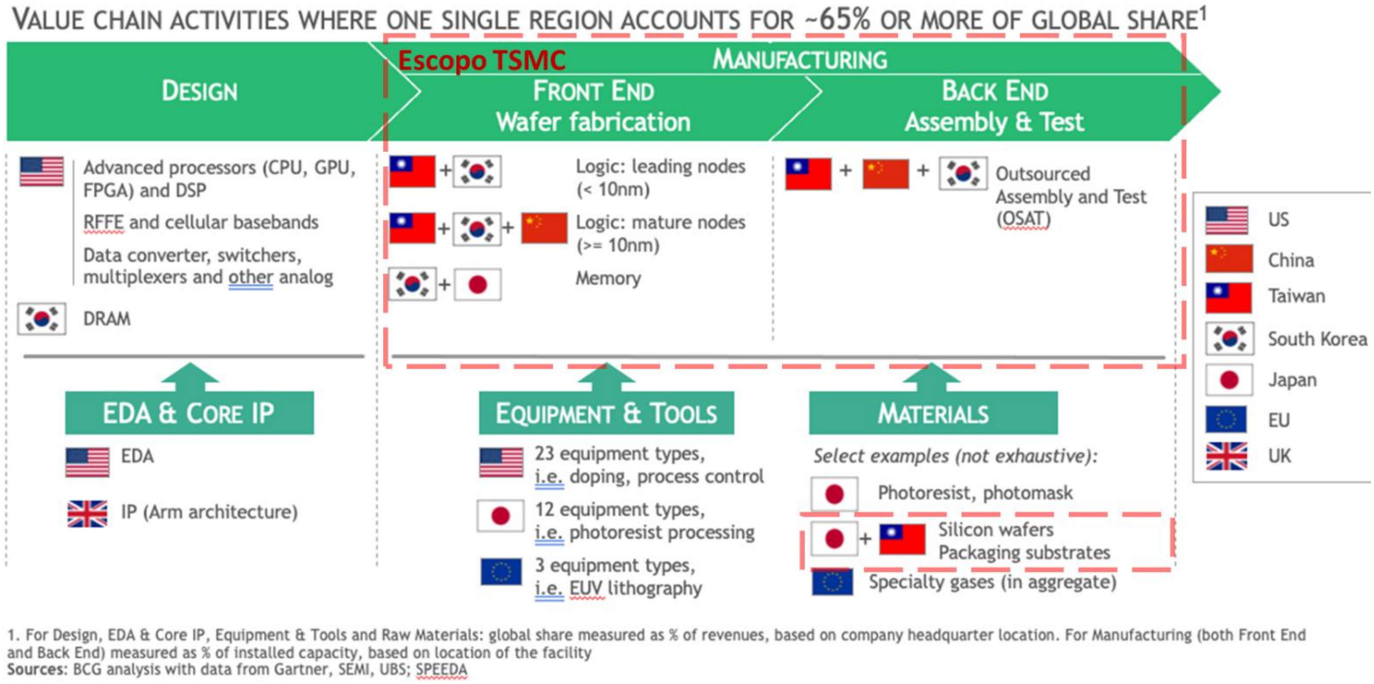

The semiconductor supply chain is broad and complex, involving many countries and stages.

https://www.semiconductors.org/strengthening-the-global-semiconductor-supply-chain-in-an-uncertain-era/

ARM

ARM, which recently completed its IPO, is focused on its intellectual property, which is essentially processor-level instruction logic and layers of complexity built on top of it, called "microarchitecture." ARM's instructions ultimately proved more efficient than those that had dominated the market until then (Intel). The instruction set logic (*Instruction Set)* is something highly complex and fundamental to how a processor works. Until then, the x86 architecture dominated the market, but with the advent of *smartphones*, ARM rose unexpectedly due to greater energy efficiency primarily. Today Apple's entire product line is designed to run on ARM processors (with ARM technology and design). Apple purchases the rights to use the microarchitecture developed by ARM on top of its instruction set, and to modify and add functionality to meet the specific needs of its products.

In the future, there is the idea that *open source* instruction sets could threaten ARM's dominance, but this is a 5 to 10-year horizon. In conversation with a market specialist working at ARM, we confirmed that the greatest existential risk for ARM is indeed such *open computing*, something analogous to *open source* for code. Many players such as Google and Meta already support these projects. RISC-V is an open-standard instruction set architecture; RISC-V is notably one of these open instructions on which no rights or *royalties* are paid. They promise to make design more democratic and accessible, and are developed by researchers in an open manner. However, we are still far from having chips comparable to an iPhone's in RISC-V.

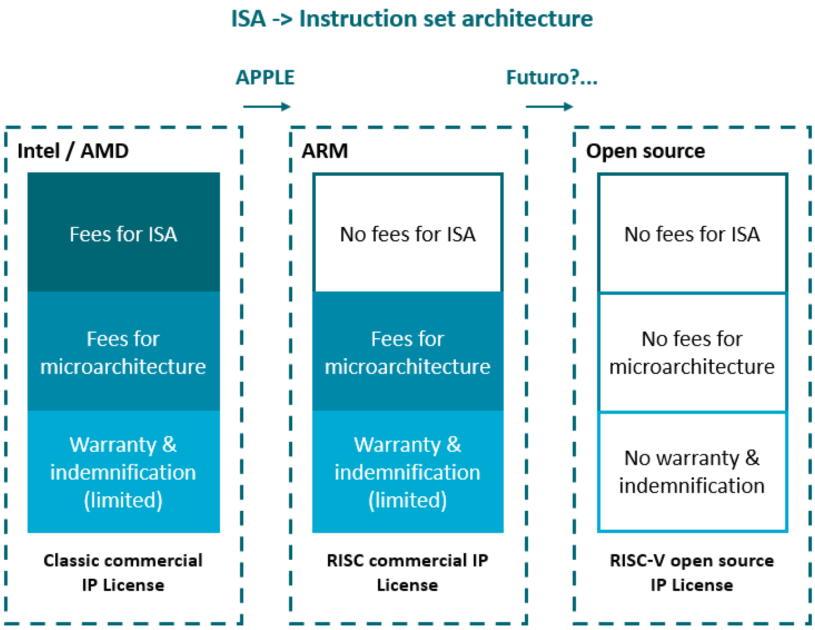

RISC-V vs. ARM vs. x86 – What's the difference? (microcontrollertips.com)

China and Taiwan: Military Tensions

"To subdue the enemy without fighting is the acme of skill. Thus the highest form of generalship is to balk the enemy's plans; the next is to prevent the junction of the enemy's forces; then to attack the enemy's army in the field; and the worst is to besiege walled cities."
A Arte da Guerra | 孫子兵法

We believe that predicting the geopolitical and macroeconomic scenario is extremely difficult. When investing in a company, our philosophy prioritizes analyzing the company's internal strengths, its market, its management, and its results. However, we understand that the main thesis risk for TSMC would be the geopolitical tension between the U.S. and China. And while we do not claim to predict a specific scenario, analyzing the likelihood of a potential conflict is important to measure the risk we are taking on the thesis, and eventually to calibrate the position size.

First of all, we understand that China does not have a very expansionist military posture. In its dispute with Taiwan, China prefers to resolve the matter peacefully, as preached by its millennia-old book still read today as a military strategy primer. Xi Jinping himself has signaled that, for better or worse, Taiwan must return to China. Furthermore, both Taiwan (ROC *Republic of China*) and mainland China (PRC *People's Republic of China)* view themselves as similar and share a common culture, which in theory would help prevent an unscrupulous attack on Taiwan.

The macroeconomic environment of mainland China is not very favorable. Starting with the deceleration of its GDP, after decades of prosperity, a large share of young people are unemployed. Its real estate sector, which accounts for approximately 20% of GDP, is facing serious difficulties. With consumer confidence at historically low levels, many projects had to be simply abandoned, becoming virtual ghost cities. China's population is likely at its peak and should shrink from here onward. Moreover, China, which already faced a food crisis in the 1960s, still bears the burden of importing large amounts of food to feed its population. Despite being a major global power, we understand that the most opportune moment for a potential invasion is not now. These macro factors, while not completely preventing it, reduce the chance of any type of invasion in the short term.

China's military power, while strong, is still far short of American capabilities. China has relatively little experience in large-scale armed conflicts in recent decades, with its last significant military incursion being the conflict with Vietnam in 1979, a long time ago. The US, in turn, has a long list of military incursions across many continents and against many adversaries. Americans have extensive experience developing and delivering new military technologies, in addition to the world's largest defense budget. Perhaps the greatest US differentiator remains its training and strategy, derived from its broadly military and interventionist policy in recent decades. As an example, we could cite that from their aircraft carriers, the Chinese had not, until 2022, taken off or landed a single aircraft. This shows a lack of confidence in that technology. From a US Naval Institute report:

"The People's Liberation Army Navy likely pursued cautious, low-intensity flight operations to maximize aircraft fuel reserves for a divert and reduce overall risk to carrier operations. Japan Maritime Self-Defense Force photos and People's Liberation Army video of the deployment showed J-15s either flew "clean," without any drag-inducing external stores or with a few relatively light air-to-air missiles. No large payloads involving air-to-surface weapons, drop tanks, or pods were noted. Japanese officials reported the Liaoning conducted more than 300 aircraft sorties during the May 2022 deployment. ***Over the 12 days of reported flight ops, that averages to fewer than 20 fighter sorties per day****, combined with dozens of helicopter flights throughout the deployment.* ***Chinese military commentators declared this to be a "decent number" of sorties for training. For some perspective on those numbers, the U.S. aircraft carrier USS Gerald R. Ford (CVN-78) recently set a record for completing 170 sorties in eight and a half hours of training."***

Lessons from the Changing Geometry of PLA Navy Carrier Ops | Proceedings - January 2023 Vol. 149/1/1,439 (usni.org)

A possible attack on Taiwan would require amphibious technology; the invasion would not be easy, as the country has few beaches and they are already heavily fortified. Moreover, the number of Taiwan's military personnel represents an expressive portion of the population, meaning Taiwan's Army can defend itself and is also well-trained. It would also have very high morale in defending against a possible invasion. Most military exercises confirm that Taiwan would have a good capacity to defend itself until Western reinforcements arrived (Japan, the Philippines, and above all the US). It is worth remembering that the US has a military presence in all allied countries in the region such as the Philippines, Japan, and South Korea, as well as a base in Guam.

On balance of our analysis, we understand that China would have a good chance of losing if the US and its allies intervened, which is very likely. And this allows us to sleep more comfortably regarding the possibility of an invasion.

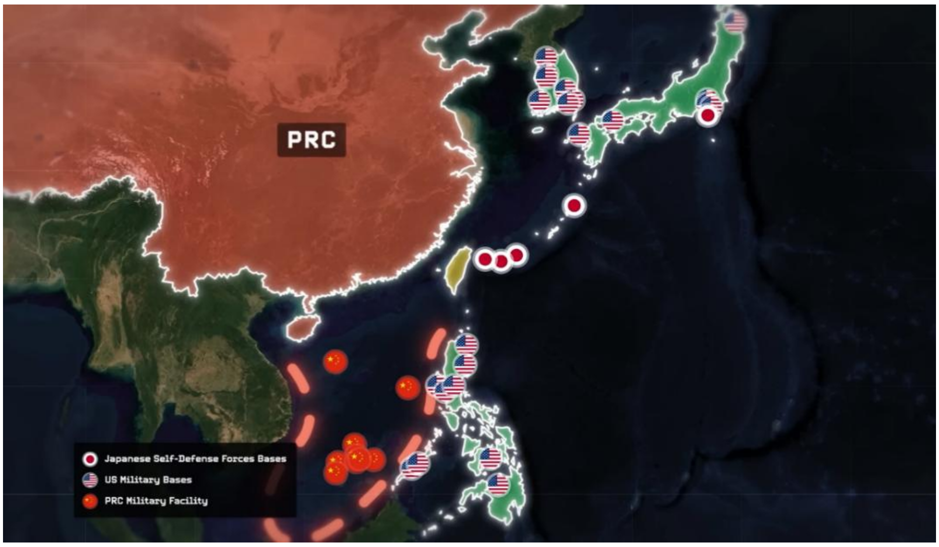

Representation of the 1st island chain: Taiwan in yellow, the Philippines below, and Japan above-right. All territories where the US already has significant military presence. [How the US & China Are Preparing to Fight Total War - RealLifeLore]

Guam (US territory) highlighted, with 8,000 troops and broad military presence, including an aircraft carrier group. In yellow, the oil route China uses that could be controlled.

American military assets in Guam.

Given the heavy dependence of many American companies (Nvidia, Apple, AMD, etc.) on Taiwanese technology, it is hard to imagine the US remaining indifferent to an invasion that could have unprecedented economic impacts for everyone.

The Cycle: Where Are We?

It is worth noting that semiconductor manufacturing is a cyclical industry, growing approximately 11% annualized over the past 40 years and responding positively and negatively to GDP and consumption. A temporary contraction of the sector following severe crises (2001, 2008) is normal, but this moment is always followed by a rapid and strong recovery in subsequent years. We are currently emerging from a downturn, and in our projection this trough should be reached in Q1 2024. This means that after that, the numbers should begin to improve on a trailing annual basis.

Analyzing inventory levels and revenue growth, we are at the best point in the cycle. The semiconductor industry is currently resuming growth and inventories are beginning to decline. The best point in the cycle to invest is now, before revenues explode and inventories normalize, which should happen in 2024 and 2025.

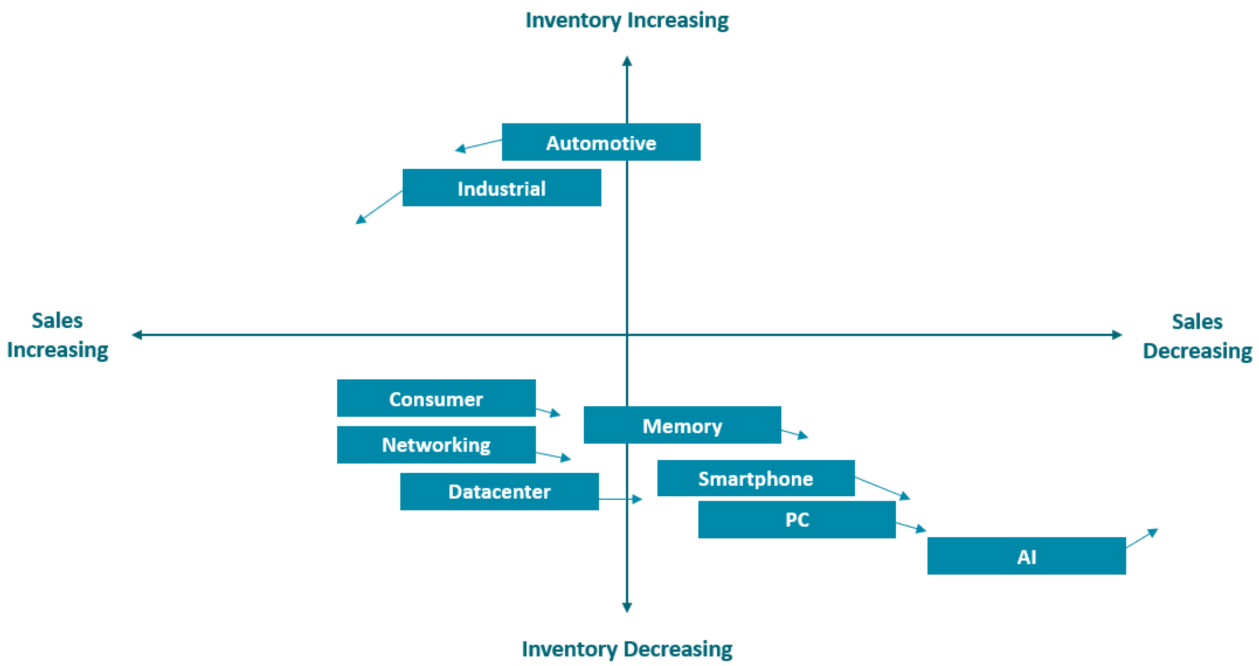

The semiconductor cycle is conducive to investment. In the coming months we are emerging from the worst point https://www.fabricatedknowledge.com/p/semiconductor-cycles-industrial-and

Artificial Intelligence and Its Implications for Semiconductors

In recent months we have seen an explosion and effervescence around *AI* and LLMs (L*arge Language Models*). The main characteristic of LLMs' operating requirements is that they naturally demand parallel processing. The natural architecture of neural networks and training algorithms, combined with the large size of recent models, push parallel computing to the extreme. The figure below illustrates how parallelism can be more efficient in processes that repeat many times.

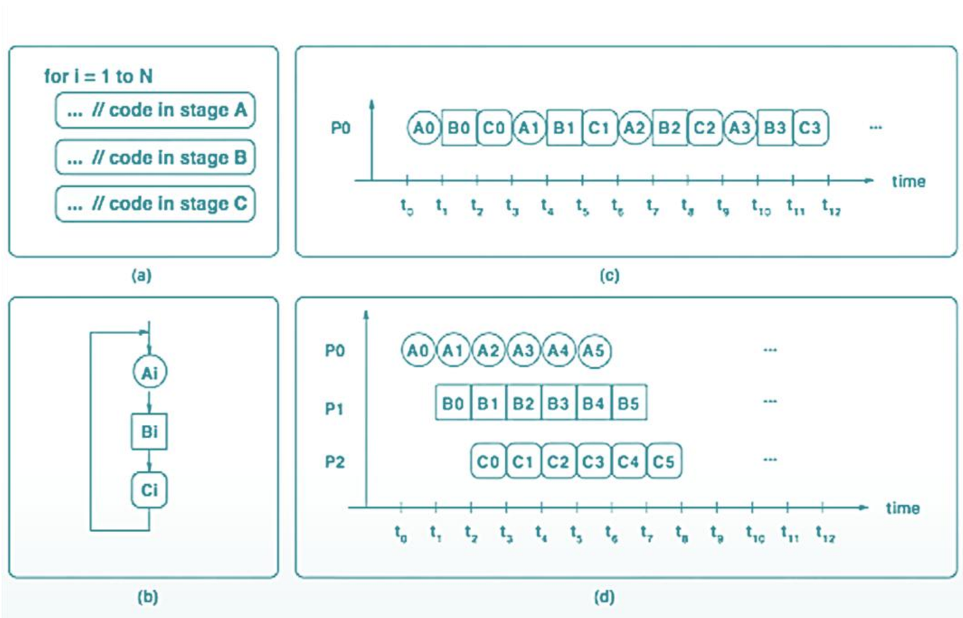

Digital Circuits - Systolic Arrays and Beyond (ETH Zurich, Spring 2017)
Top: serial processing; Bottom: parallel processing

Artificial Intelligence: Parallelism

In conventional computing, heavy use of parallelism was not necessary, but with the advent of current Language Models, demand for this type of computing exploded. Very opportunely, however, Nvidia already had decades of advances and specific knowledge in parallel processing.

CUDA is a library developed in 2007 to facilitate the use of parallelism on Nvidia's specific hardware at the time. The initial application was 3D graphics processing. Its cards already used a great deal of parallelism for graphics processing: post-processing, real-time physics simulation, texture processing, and more.

Although the initial problem was predominantly in visual processing, other industries soon began discovering and using CUDA technology for many computationally intensive and repetitive tasks; subsequent uses included scientific research, medical imaging, and modeling of the most varied types (oil and gas and even econometric).

Therefore, when LLMs came to the fore, Nvidia had the resources and skills at hand, having already built and matured its software and hardware to solve the problem in question: efficient parallelism. This capability has been very well rewarded recently, in the context of the explosive demand for its A100 and H100 cards, which are at the eye of the revolution that LLMs have brought in recent months. The images generated by *AI* in this document were most likely processed in a Microsoft data center using Nvidia cards, which in turn used chips manufactured by TSMC at 7nm.

Therefore the benefit here is direct, since Nvidia, despite controlling the entire *Stack* technology stack, does not manufacture its own 7nm chips for the cards it sells. And TSMC dependence in this product line is 100%. It is estimated that demand for *AI* chips for TSMC should grow at a compound rate of 50% over the next 5 years — a market that should multiply 7x in 5 years. And we believe that demand for *AI* is only at the beginning.

Artificial Intelligence: Massive Capital Investment

In the AI era, it is increasingly evident that Capital is becoming an ever-greater barrier. In conversations with industry insiders, it was reported that an investment in the range of US$1bn in data centers is already considered small by today's standards.

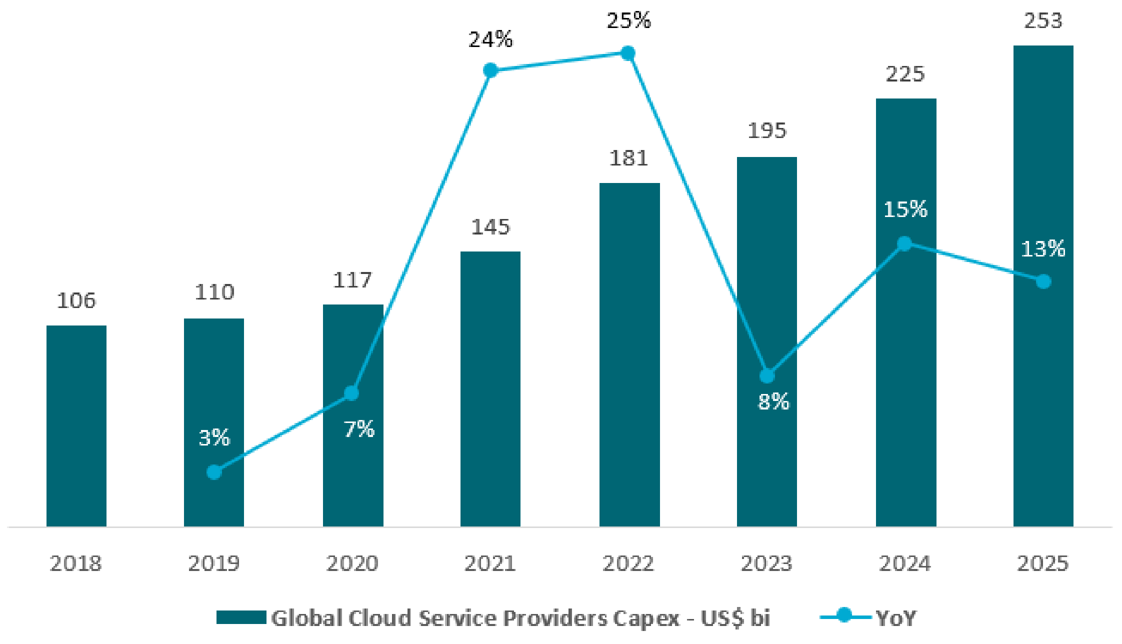

Counterpoint Research

Also in these conversations we learned that *hardware* hardware also tends to depreciate more quickly, largely due to obsolescence. Previously, a useful life of around 6 years for equipment and servers was considered reasonable; today this is trending toward depreciation in as little as 4 to 5 years. This means that operating costs are becoming increasingly prohibitive, requiring ever more capital (capex and opex in the form of depreciation) merely to remain in the game. These requirements end up "drowning" smaller players. In poker, it is as if a player with few chips were gradually "bleeding out" to *blinds* increasingly larger blinds. This logic applies both to cloud players that invest hundreds of billions of dollars in data centers and to TSMC itself, which in turn also has to invest billions in R&D, operations, and machinery. In semiconductors, the *blind** is already prohibitively high.

In *AI*, all of this prohibitive capex logic becomes even more intense and below we explain this with some numbers. *AI* is represented in the "GPU server" column because basically the *workload* of *AI* workloads do not run on conventional CPUs. Until a few years ago, demand was much more for CPU, but with *AI* and LLMs there is a massive migration to GPUs. Looking at the numbers, it is clear that *Capital Costs* (related to capex and depreciation) are comparatively much larger than other maintenance costs (cooling, upkeep, etc.).

Blind: minimum bet to enter a round of Poker.

	CPU Server	GPU Server (AI)
Server Capital Costs		
Upfront Server Capex	$15.000	$350.000
Useful life in years	6	6
Cost of Capital	13%	13%
Server Capital Costs	$301	**$7.026**
Server hosting Costs		
Electricity Cost	$0,09	$0,09
Utilization rate	80%	80%
Power Usage Effectiveness (PUE)	1,25	1,25
Electricity Cost per kW per month	$63,5	$63,5
Colocation cost (i.e. rack sapce rental cost, infra, insta	$120,0	$120,0
Total Server Hosting cost per month	**$183,5**	**$183,5**
Server Power Consumption	$1.200	**$10.200**
Total Server Hosting cost per month	$220	**$1.872**
Total Server Cost per Month	**521,3**	**8.897,7**
Capital Costs / Total Costs (per month)	58%	**79%**
		Capital é ainda mais relevante !

Semianalysis Research

This benefits TSMC in that saving on GPUs is becoming increasingly important (increasing demand for TSMC's services) and because TSMC's activity is tied to an ever-larger portion of the costs of operating a data center (which transfers bargaining power to TSMC and its advanced technologies). The *Big Techs* increasingly try to avoid the so-called "*Nvidia tax*": an allusion to the "tax charged by Nvidia to operate in *AI*.

ASICs and *Custom Chips*: The Entire *AI* AI Market in TSMC's Hands

As a result, all relevant *AI* players are creating their customized chip versions for *AI*. Nvidia is indeed a technological leader in GPUs and *design* card design has been in its DNA from the beginning. Therefore, the *Big Techs* social media (Meta) and *advertising* (Google) had to step a little outside their comfort zone and flex muscles they were less accustomed to — that is, they began directing efforts toward creating their own semiconductors. But the *gap* is closing. Google was the first to venture into this type of project. Today all *Big Techs* that take *AI* seriously have their own semiconductor design initiative, called "*custom sillicon*". And all of these projects have one thing in common: they are ordered from TSMC.

For this task, these projects are increasingly moving in the direction of **ASICs, acronym for "*Application-specific integrated circuit"*,** which are circuits specific to a task and which provide better *unit economics* in the long run. Basically there is a *trade off* expected between functionality and efficiency; by making these cards more specific (for LLMs and neural networks), a better final operating cost is obtained. ASICs are chips designed from scratch to meet a specific demand, in this case running neural networks. Their entire design and architecture is optimized for this purpose. On the other hand, they sacrifice flexibility by having a very specific purpose.

It is worth reiterating that the design cost alone of an ASIC for *AI* runs in the range of at least tens of millions of dollars and requires access to highly specific and qualified labor. The larger the order (with TSMC), the lower the unit cost achievable in negotiations.

Google began its TPU (*Tensor Processing Unit*) ASIC project in 2013, with the first versions delivered in 2015. Today it is heading toward its 5th iteration and in "*code red*" mode after OpenAI suddenly took the AI world by storm. All its *tensor chips* are manufactured at 5nm with TSMC.

Amazon announced plans to produce its AI ASIC chip in 2018. Today it is already running a portion of its LLM demand. Called *Trainium* and *Inferentia* are already in their 2nd iteration. In conversations we had with Amazon specialists, significant savings in capex and operating costs were reported when running LLMs. Amazon also has a partnership with Anthropic (a competitor to OpenAI), which develops its language models from scratch to run on Amazon's cloud: "*We are working closely with AWS to develop our future foundation models using Trainium chips*" said Tom Brown, co-founder of *Anthropic*. "*Trainium2 will help us build and train models at a very large scale, and we expect it to be at least 4x faster than first generation Trainium chips for some of our key workloads. Our collaboration with AWS will help organizations of all sizes unlock new possibilities, as they use Anthropic's state-of-the-art AI systems together with AWS's secure, reliable cloud technology.*". All *Trainium e Inferentia* cards are manufactured at 7nm with TSMC.

Additionally, we have Microsoft's efforts, having announced its ASIC *AI* chip called "Maia." It is the first complete server processor with liquid cooling built by Microsoft. Maia chips will be manufactured at 5nm with TSMC.

Lastly, Meta is also designing an ASIC for *AI*. Its chip called MTIA (*Meta Training and Inference Accelerator*). MTIA chips will be manufactured at 7nm by TSMC.

Therefore, as we have noted, all future demand for *AI*, via Nvidia or other *Big Techs* is directed into TSMC's hands. The intelligence that will revolutionize the 21st century will be forged/cast in Taiwan, by TSMC.

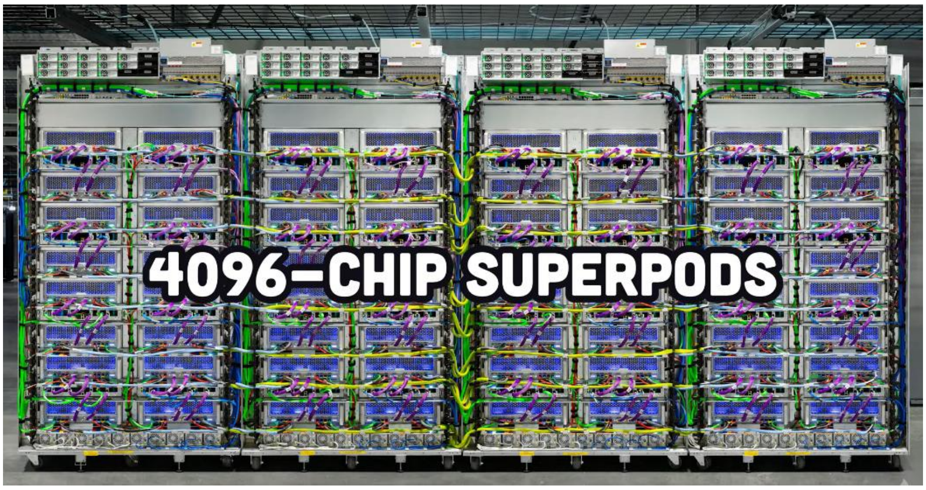

Google: TPU v5, superpods of 4,096 chips with optical network; 5nm TSMC

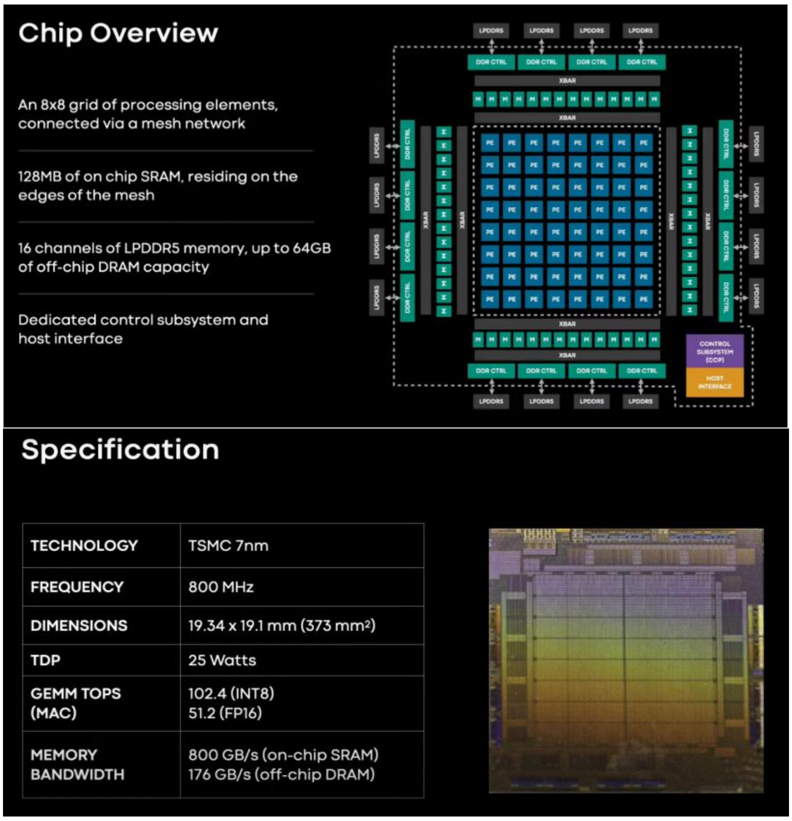

Meta: MTIA chips; 7nm TSMC

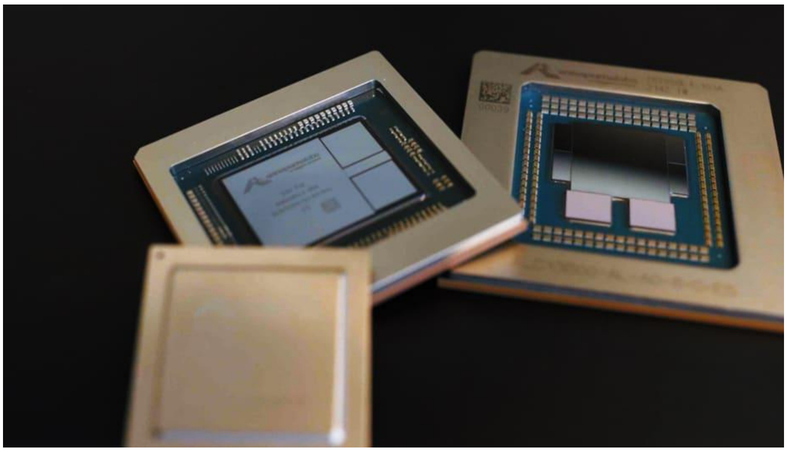

Amazon: From left to right, Inferentia, Trainium and Graviton, from Seattle on July 13, 2023. 7nm TSMC

Microsoft: Maia chips; 5nm TSMC.

The diagram below simplifies the view of the *trade off* between flexibility vs. functionality and efficiency. With the maturation of LLM and AI technologies, the *Big Techs* will likely chase the *bps* of efficiency, with hardware designed for its specific purpose.

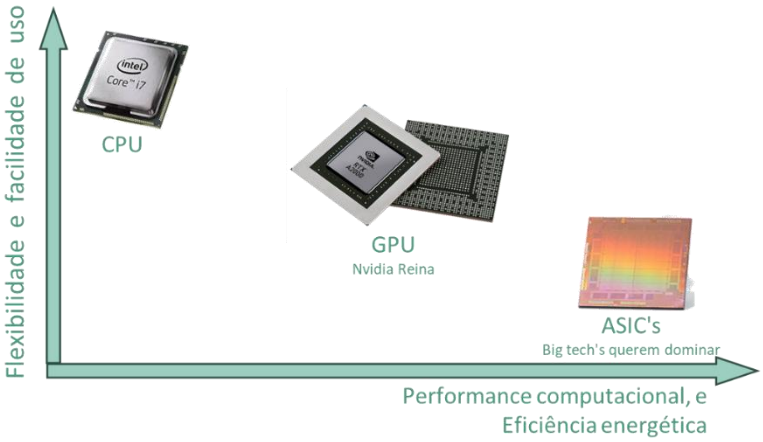

CPU, GPU, ASICs: trade-off between flexibility and performance.

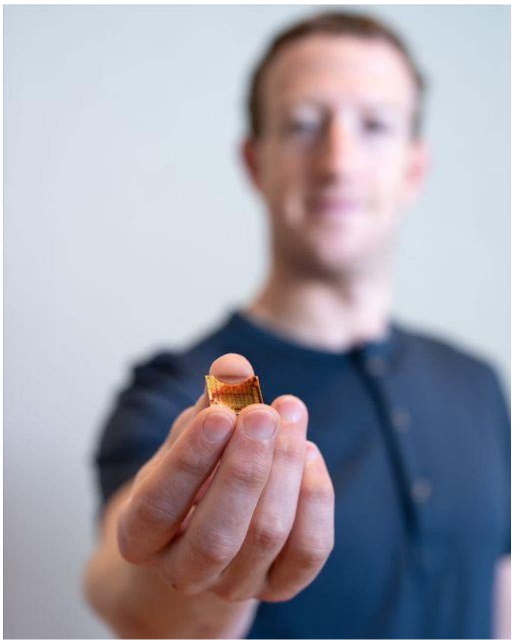

Zuckerberg announcing the new ASIC chips for AI (MTIA)

TSMC Is Also a Services Company, It Sells the Shovel, Not the Gold

Microchip is a "*pick and shovels*". As it is far too risky for customers to engage in the activity of *design* and product management, the seller *pick and shovels* (TSMC) benefits from a large market, where the greater risk is borne by TSMC's customers. Beyond being a supplier of shovels (selling basic technological inputs), TSMC also offers scale, technology, and *know-how*, which its customers alone would never achieve. Therefore, it is even better than selling shovels in a gold rush, given TSMC's unique character.

When the customer places an order for a batch for 2026, the one bearing the greatest risk is the customer itself — AMD, Nvidia, Xilinx, etc. If their technology does not succeed (due to better competition), their product goes out of fashion, or demand simply declines (gaming cards post-pandemic), TSMC already has a contractually predetermined price for services and a minimum quantity from the initial batches. TSMC demonstrates high predictability in its revenues, establishing rigorous contracts that minimize the risk of default by customers with respect to payment for services rendered. On the other hand, if the customer manages to “find the gold” and achieve great success, the merits also remain with the customer at the final end. Nvidia is surfing such a moment, where its GPU cards for *AI* found an excellent market and is being rewarded for its product with a 70% margin at this moment. But if for any reason its product insertion had failed miserably, they would still be paying TSMC exactly as stipulated in the contract signed 2 years ago.

Thus the scenario is riskier in the long run for TSMC’s customers than for TSMC itself. TSMC is in its essence a services company, not a products company. It develops technologies and processes to serve its customers’ demands first and foremost.

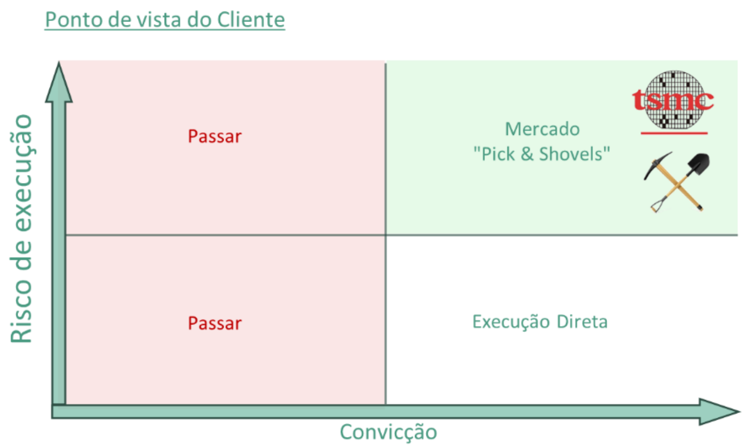

TSMC bears less risk than its customers. It sells a basic input to those seeking the gold (Apple, Nvidia, etc.)

Management, Culture and Taiwan

We understand that TSMC has management, culture, and governance of the highest quality. We see stability in management, and the group of 25 directors and vice-presidents has an average of 18 years at the company, of which 14 have 25 or more years at TSMC. Looking at their backgrounds, the vast majority obtained *PhD* degrees from leading universities and accumulated, even before TSMC, years of experience in the semiconductor industry. TSMC’s operational execution is itself an indication of high quality in capital allocation and management. Morris Chang, TSMC’s founder, is seen as a legend in the semiconductor industry, being responsible for the concept of the *foundry* dedicated foundry. At the same time, he recognized that developing this industry was fundamental to Taiwan’s geopolitical affairs. In his country he is seen as a hero.

Taiwan offers employees who culturally have strong adherence to the strategic importance of TSMC for the country. The *fit* cultural fit between the country and TSMC’s activities is high, and the cost of living in Taiwan is approximately half that of the US, allowing TSMC to have lower labor costs. Replicating Taiwan’s structure in other countries could be very difficult at the same margin levels. The US has an interest in replicating this in other countries to reduce dependence on Taiwan, which is negative for the thesis, as they are in truth seeking an alternative to TSMC. On the other hand, we believe the Taiwanese government knows that one of the reasons the US defends Taiwan is precisely the irreplaceability of the microchip supply chain based in its country.

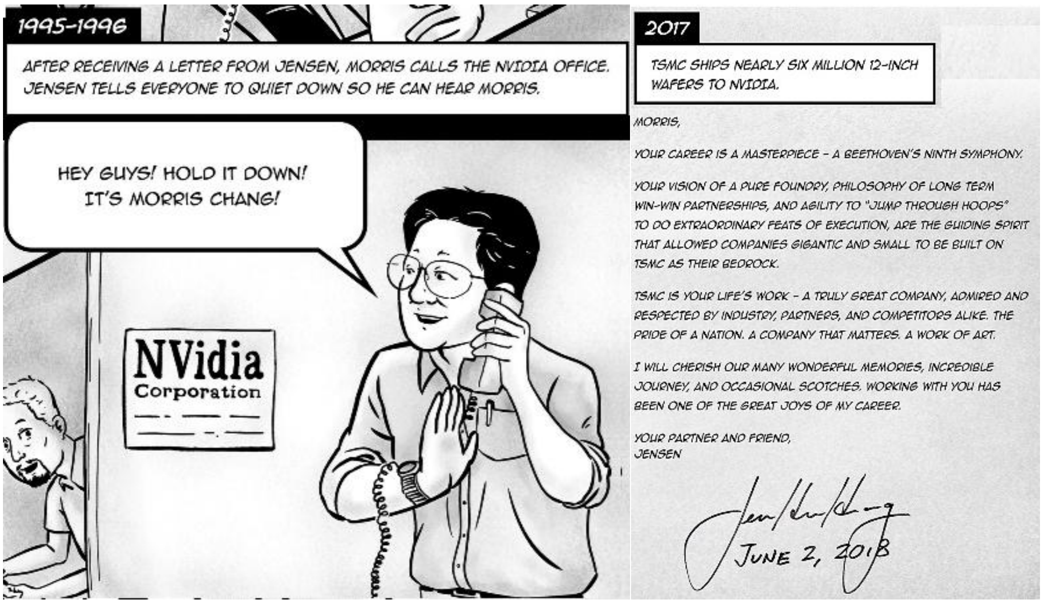

Comic depicting the relationship between Morris Chang and Jensen Huang (CEO of Nvidia)
https://blogs.nvidia.com.tw/2018/06/28/nvidia-ceo-talk-about-friendship-with-tsmc-ceo-morris-chang/#

From left to right: Tim Cook (Apple), C. Wei (CEO of TSMC), Mark Liu (Chairman of TSMC), Morris Chang (founder)

Value Creation

As an investment, we believe that investing in a value-generating company with *Moats* is the path to good returns. When analyzing TSMC's numbers, many things stand out positively: high profitability, neutral working capital, and strong cash generation. But perhaps the most impressive has been the company's 30 years of exponential growth sustaining *EVA (economic value added)* margins at a very high level of 20% to 30%.

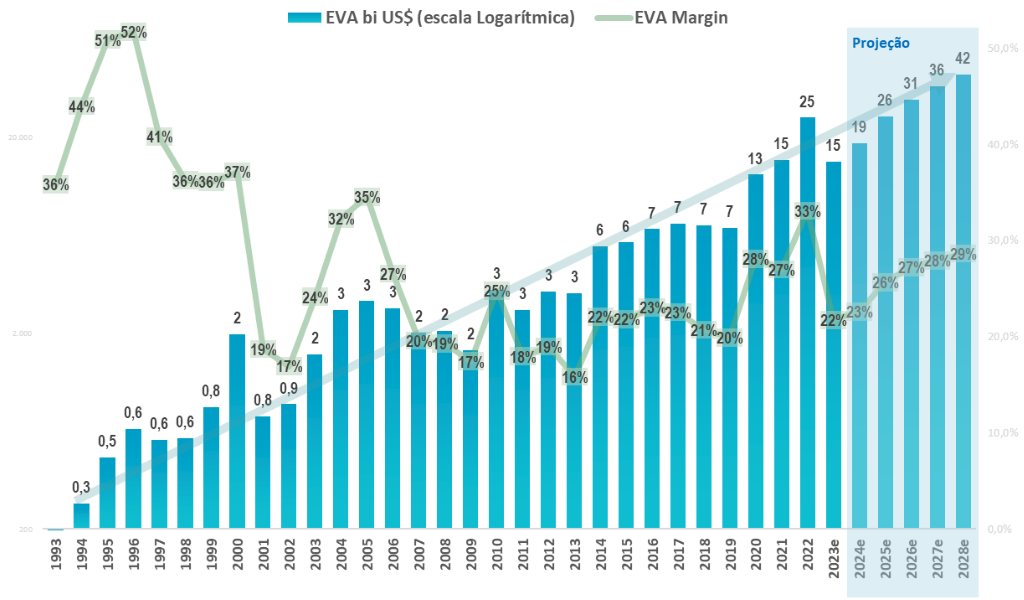

Value generation, on a logarithmic scale!

To put this in perspective, below we have a comparison between TSMC's average EVA margin over the past 5 years against the averages of US companies. Only the top 10% (90th percentile) of companies in this metric have an EVA margin of 9.7% or more, while TSMC delivers an impressive 25%!!

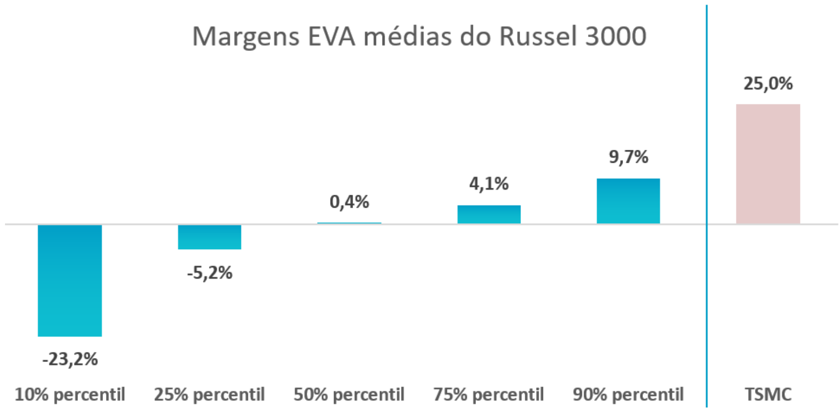

Value generation data for Russell 3000 companies

Conclusion

We believe TSMC is a fantastic company, with a history full of successes and numbers to corroborate them. We see high quality in its management, market of operation, and corporate structure. It is a company with strong *Moats* moats of scale and differentiated process with its advanced technology. Its value creation has no parallel in any industry and is a major asset in the business valuation. Its track record, unblemished reputation, and strong customer-focused culture make the investment even more interesting. We understand that the geopolitical issue with China is the major risk factor, but weighing everything and with a well-sized position in the portfolio, we see positive asymmetry for the investment.

We thank you for your trust,

SFA Investimentos Team

Manager
SFA INVESTIMENTOS LTDA
Rua Gomes de Carvalho, 768 – 8º andar
Vila Olímpia – CEP: 04547-003
São Paulo - Brasil
Tel: +55 11 2780-0690
ri@sfainvestimentos.com.br

Administrator
BNY MELLON SERVIÇOS FINANCEIROS DISTRIBUIDORA DE TÍTULOS E VALORES MOBILIÁRIOS S.A.
Av. Presidente Wilson, nº 231 – 11º andar
Centro – CEP: 20030-905
Rio de Janeiro - Brasil
Tel: +55 21 3219 2500

Custody
BNY MELLON BANCO S.A.
Av. Presidente Wilson, nº 231 – 11º andar
Centro – CEP: 20030-905
Rio de Janeiro - Brasil
Tel: +55 21 3219 2500
www.bnymellon.com.br