Porter & Co.

THE BIG SECRET ON WALL STREET

The Parallel Processing Revolution

- How The "Big Bang That No One Noticed"...
- Created A Brand-New, Rapidly Expanding Universe

FROM THE DESK OF PORTER STANSBERRY

The Parallel Processing Revolution

How The "Big Bang That No One Noticed"...
Created A Brand-New, Rapidly Expanding Universe

The "Big Bang" is still going on.

The mysterious explosion at the dawn of time continues to reverberate outwards... in a chain reaction that's actively continuing today.

As theoretical physicists will tell you... ever since our universe first spawned from a tiny pinpoint of energy, it's just kept getting bigger. At an *accelerating rate*, no less. According to one standard measure called the Hubble constant, right now the universe's many pieces are speeding outwards into space at the rate of about 45 miles per second.

We don't know the source of the powerful force pushing the stars and galaxies apart, but Albert Einstein gave it the Greek name "lambda," representing a cosmological constant. Scientists today call it **"dark energy."** Whatever it is, it's far stronger than gravity – and capable of ripping apart the very fabric of space.

The only reason dark energy hasn't pulled the universe into smithereens (yet...) is because of a second mysterious, stronger-than-gravity force that holds the cosmos together amid the expansion. Scientists call this unknown force **"dark matter."**

That's a lot of "unknown." And a whole lot of cosmic-level power.

And, fascinatingly, it mirrors the forces at work in the Parallel Processing Revolution... a rapidly expanding industrial universe, kickstarted by what we call **"The Big Bang That No One Noticed."**

Computer-chip maker Nvidia quietly pioneered a lightning-fast graphics chip in 2006. Then in 2012, three young researchers from the University of Toronto proved the chip could be used in all computers, not just for graphics. Since then, the parallel-processing industry has exploded from a tiny, nerdy niche to a $6 billion global industry that's slated to reach $1 trillion by 2030.

The tiny chips (and their descendants) created by Nvidia now power nearly every piece of technology we own, from smartphones and laptops to medical equipment and cars. The parallel-processing revolution has sparked an industry sea change on par with the Gutenberg printing press... the Industrial Revolution... and the internet itself.

In that way, the birth of the semiconductor industry is truly a "Big Bang"... one that spreads outward with tremendous, accelerating force... and one that's held together during its expansion by a powerful gravitational pull...

The Parallel Processing Revolution has its own "dark energy" – the vital fuel sources that power the exponential growth of the industry. It even has its own "dark matter" – infrastructure that keeps the complex and ever-evolving semiconductor sector from collapsing.

In this compendium of reports, we'll take a deep dive into the extraordinary universe of the Parallel Processing Revolution.

- First, we'll zoom in and look at the brilliant core where it all begins... the source of the original "Big Bang," Nvidia.
- Then, we'll examine the universe's elemental building blocks... three little-known companies we call "Nvidia's Silent Partners."
- We'll take a closer look at three energy titans that propel the industry...
- And, finally, we'll reveal a key infrastructure company that holds it all together.

We hope you enjoy this in-depth investigation into the Parallel Processing Revolution. Please note that these companies are not official recommendations, but at Porter & Co., we believe they will present some of the most valuable buying opportunities of the next decade. We suggest you study them closely.

To your success,

Porter Stansberry

Porter & Co.

Table of Contents

PART 1: HOW IT ALL BEGAN

It all started with Nvidia. Parallel computing has revolutionized the tech world, and **Nvidia (Nasdaq: NVDA)** stands at the forefront of this transformation. Founded in 1993 when three humble computer engineers sat down in a Denny's and drew up plans for a graphics-focused startup, today Nvidia is the single biggest player in the parallel processing revolution – and the world's third-most valuable publicly traded company, behind Apple and Microsoft.

Nvidia's groundbreaking CUDA (Compute Unified Device Architecture) platform opened up the parallel-processing power of GPUs (graphics processing units) for a wide range of applications beyond graphics. CUDA was soon adopted by tech giants like Google, Facebook, and Baidu, which operates China's largest search engine. And from that point on, there was no stopping Nvidia.

Over much of the past decade, strategic acquisitions and product developments have solidified Nvidia's leadership in parallel computing – and catapulted it to a nearly $3 trillion company with a chokehold on the semiconductor industry.

Remarkably, it all happened almost by stealth... making this **"The Big Bang That No One Noticed."**

In this report, you'll find a full analysis of Nvidia's business... why the company is the undisputed center of the parallel-processing universe... and why it simply cannot be unseated.

Porter & Co.

THE BIG SECRET ON WALL STREET

The Big Bang That No One Noticed

- The Parallel-Processing Revolution Has Only Just Begun
- Trying to Speed Up Video Games, Nvidia Rocked the Tech World

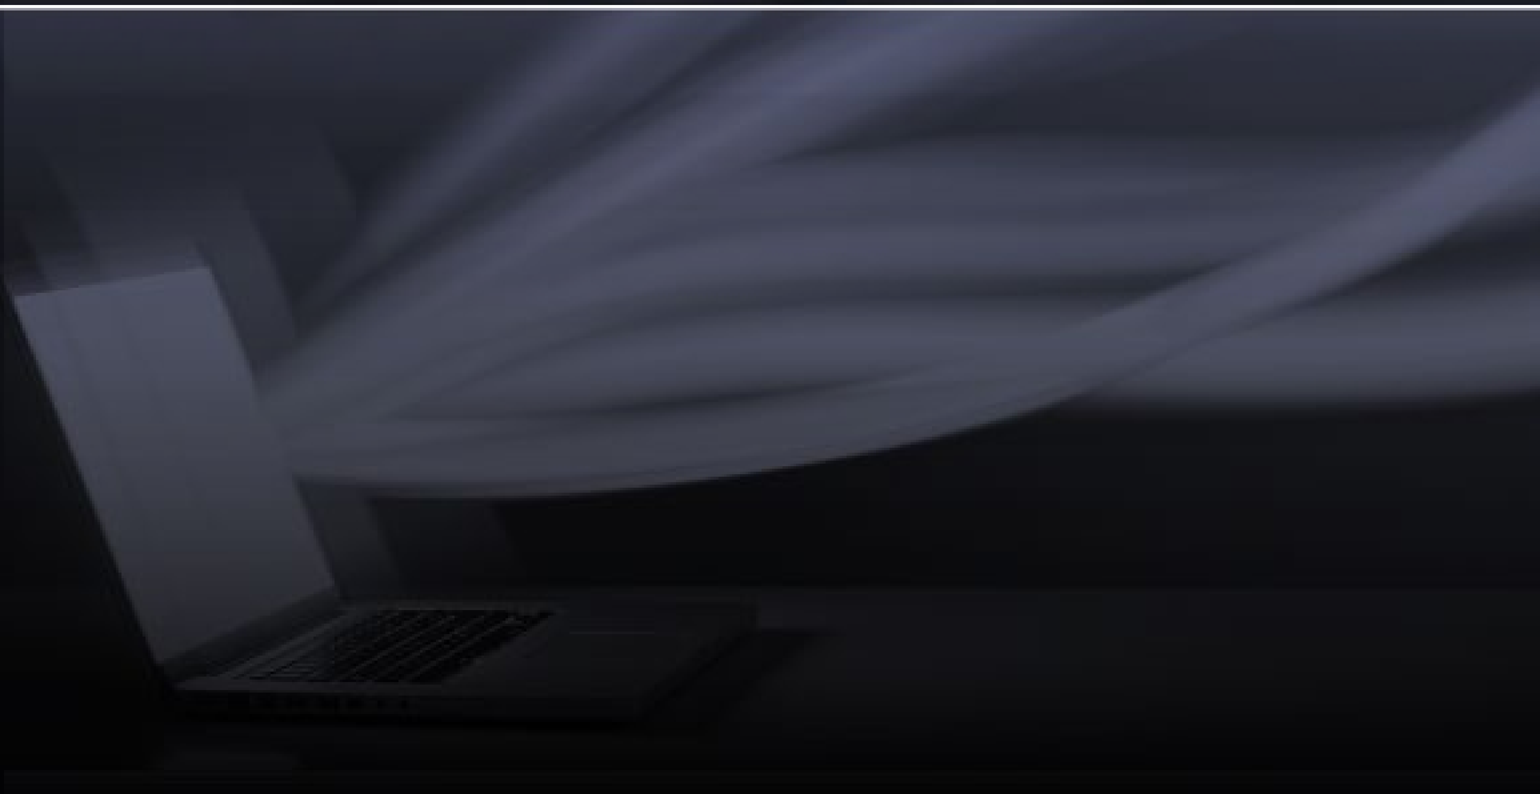

FROM THE DESK OF PORTER STANSBERRY

The Big Bang That No One Noticed

The Parallel-Processing Revolution Has Only Just Begun

Trying to Speed Up Video Games, Nvidia Rocked the Tech World

In October 2012, three obscure academics at the University of Toronto accidentally changed the world.

Earlier that year, in search of little more than bragging rights, PhD candidates Alex Krizhevsky and Ilya Sutskever, and their advisor, Professor Geoffrey Hinton, entered the ImageNet Large Scale Visual Recognition Challenge ("ILSVRC"). This annual competition aimed to advance the field of computer vision – training computers to recognize that a photo of a man or a woman represents a person, and that an image of a poodle or a German shepherd depicts a dog.

The competition, which that year took place in Florence, Italy, centered on the ImageNet dataset, a collection of 14 million individually-labeled images of everyday items across thousands of categories.

Participants in the Challenge set out to design algorithms – an algorithm is a set of rules that programmers create – to enable computers to correctly (and autonomously) identify as many of these images as possible. This feat, known as object categorization, was already recognized as one of the most fundamental capabilities of both human and machine vision, and was an early goal of the nascent field of machine learning.

The University of Toronto team's algorithm – dubbed AlexNet, after its lead developer – won the challenge that year.

In fact, it trounced the competition. AlexNet performed significantly better than any algorithm ever had, by a wide margin. It identified images with an error rate of just 16%, while previous-year winners had error rates – that is, the pace at which algorithms incorrectly identified images – of 25% or more.

In the field of computer vision, this margin of victory was akin to Roger Bannister running the first sub-four-minute mile in 1954 – a feat that runners had been trying (and failing) to accomplish for nearly a century. It was a huge improvement over the small, incremental progress that had previously been made to date, and was a major step in the evolution of the industry.

However, while the team's low margin of error was extraordinary, what was particularly noteworthy was *how they achieved it.*

In short, Alex and his PhD colleagues had trained their algorithm on a graphics processing unit ("GPU") – a specialized computer processor originally designed to speed up graphics rendering in video games – rather than the standard central processing units ("CPU") that run traditional computers.

How the Shift to GPUs Changed Everything

Prior to AlexNet in 2012, ILSVRC teams had trained their algorithms exclusively on CPUs. CPUs are the brain of a computer. They're fast and powerful, but they have a significant limitation: they can only execute one operation or instruction at a time, one after another – a process known as serial computing.

In serial computing, the speed of a CPU is determined in part by the number and density of transistors built into its circuits. (Transistors regulate electrical signals and are the basic building blocks of modern computers.) All things equal, the greater the density of transistors a CPU has, the higher its "clock speed" – the number of operation cycles it can carry out per second – and the greater its processing power.

In the 1970s, Gordon Moore, the co-founder of Intel, noted that the number of transistors that can be manufactured into a microprocessor doubles every 18 to 24 months with minimal increase in cost. This observation became known as Moore's Law and helped the industry anticipate that CPU performance would improve at roughly the same rate, or around 50% per year.

However, with serial computing, even the most advanced supercomputers weren't powerful enough to efficiently run neural networks like AlexNet – machine learning programs that make decisions in a similar manner to the human brain. For example, the number of computational operations required to train just one advanced algorithm can rival the total grains of sand on earth… that is, it's an unfathomably large number of calculations that even the most powerful CPU would struggle to execute.

GPUs were originally designed to process graphics, which required handling hundreds or thousands of individual pixels – short for "picture elements," the smallest units of a digital image or display. Modern displays can contain 8 million or more individual pixels. Because of this, GPUs must be able to execute multiple independent operations simultaneously.

This process is known as parallel computing. A GPU can carry out tens of thousands of operations at once, creating a total processing capacity that is exponentially greater than that of a CPU. And as shown in the chart below, the relative performance advantage of GPUs versus CPUs continues to increase over time:

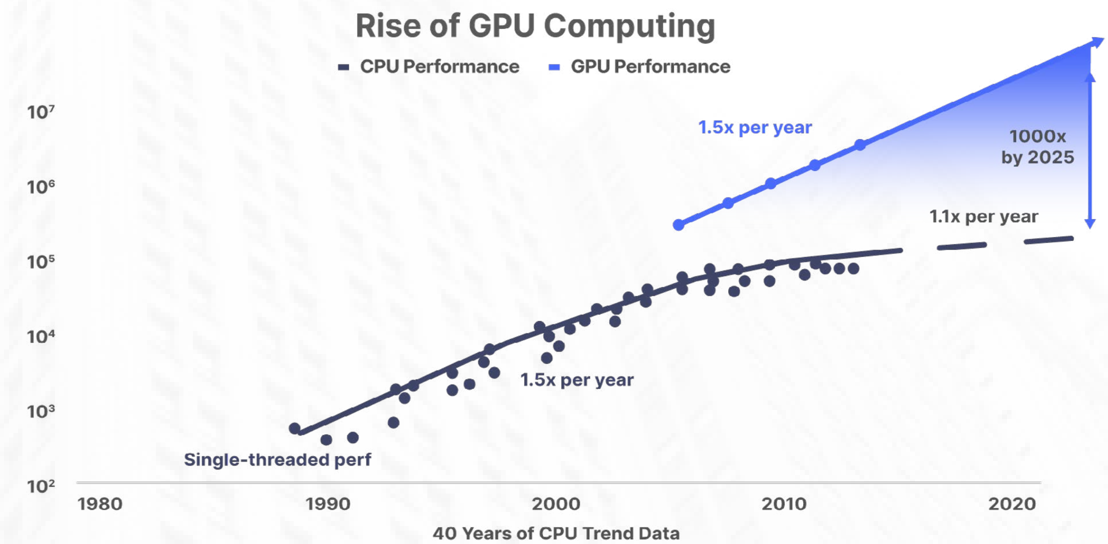

Source: Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotum, L. Hammond, and C. Batten New plot and data collected for 2010-2015 by K. Rupp

(This growing advantage is the result of a significant slowdown in the pace of CPU performance increases to around 1.1x per year – versus the 1.5x per year Moore's Law predicts – as increases in transistor density are beginning to run into the limits of physics. Meanwhile GPU performance has continued to increase at that same 1.5x per year rate.)

A simple analogy can help to explain the differences between CPUs and GPUs: A CPU is like the owner of a burger joint that serves hundreds of customers a day. The owner could potentially make all the burgers himself – a simple but time-intensive task – but it would leave no bandwidth to manage other aspects of the business. Instead, the owner could hire line cooks to make the burgers for him. In this case, a GPU is like a specialized line cook with 10 arms that can make dozens of burgers at the same time.

By running their algorithm on GPUs rather than CPUs, the AlexNet team was able to dramatically outperform other challengers.

The University of Toronto team's victory was a big deal in the computer-science world, but at the time it didn't raise any eyebrows in the broader technology universe. More than a decade later, though, it is remembered as the Big Bang moment for the artificial-intelligence ("AI") and machine-learning revolution that is sweeping the world today.

And one company has been leading the way...

How Nvidia Became the King of Parallel Computing

Parallel computing has revolutionized the tech world, and Nvidia (NVDA) stands at the forefront of this transformation.

Founded in 1993 by Jensen Huang, Chris Malachowsky, and Curtis Priem, Nvidia set out to create the first 3D graphics cards for consumer PCs. At the time, consumer video graphics were flat (two-dimensional). High-end video graphics required expensive professional workstations that were primarily the domain of the military and big-budget movie studios (*Jurassic Park* would enthrall movie-goers with its computer-generated imagery of dinosaurs that same year).

The company's first graphics cards were a big hit with consumers… so much so that it soon faced competition from as many as 90 other companies producing similar cards.

Nvidia's first major GPU breakthrough came with the RIVA 128 in 1997. The company used emulation technology – essentially using software to test its processors in a virtual rather than real-world environment – to speed up the development of this processor. This new process allowed Nvidia to begin bringing its new GPUs to market in just six to nine months versus the industry standard of 18 to 24 months. This is a key factor that helped establish Nvidia's lead in this field, outpacing potential rivals.

Building on the success of the RIVA 128, the fast-growing California-based company released the GeForce 256 in 1999. This GPU further differentiated Nvidia by introducing programmable shaders, a feature that enables developers to create more realistic levels of light, darkness, and color by leveraging parallel processing capabilities to the fullest.

In 2006, Nvidia unveiled the CUDA (Compute Unified Device Architecture) platform. CUDA is a proprietary framework for general-purpose GPU computing, which opened up the parallel processing power of GPUs for a wide range of applications beyond graphics for the first time. Putting all these advancements together, CEO Jensen Huang envisioned a "full-stack" solution, equipping developers across industries with all the tools they needed to tap into parallel processing power.

CUDA's impact was profound but not immediate. At the time, there simply wasn't a market for general-purpose GPU computing. But it was this CUDA platform that enabled the University of Toronto researchers to use Nvidia GPUs for their contest-winning AlexNet algorithm some six years later, igniting the use of GPUs in AI and machine learning. And it helped further distance Nvidia's offerings from those of its competitors.

This CUDA breakthrough soon led to widespread adoption by tech giants like Google, Facebook, and Baidu, which operates China's largest search engine. And over the past decade, strategic acquisitions and product developments have solidified Nvidia's leadership in parallel computing.

For instance, in 2020, Nvidia acquired networking products company Mellanox, which helped Nvidia enhance data transfer speeds for its high-performance computing ("HPC") data centers. And in 2022, Nvidia introduced its new Hopper GPU architecture – designed specifically for modern data centers, AI, and HPC use – as well as its Grace CPU, its first-ever data center CPU. (These two products were named after Rear Admiral Grace Hopper, one of the first female computer scientists and a pioneer of computer programming.)

Then in 2022, the rise of generative AI models cemented Nvidia's dominance in HPC. Using Nvidia's advancements, OpenAI released ChatGPT in November 2022, creating a watershed moment that showcased practical applications of AI for the first time.

Some of the real-world use cases for ChatGPT and other generative AI applications include customer service (providing automated and multilingual support), content creation (generating high-quality, human-like text for websites, blogs, social media, and marketing materials), marketing (providing personalized product recommendations to customers), legal and compliance (quickly analyzing legal documents, extracting relevant information, and summarizing them), and real-time language translation.

Nvidia's innovations in parallel computing were instrumental to these advancements. Its GPUs, with their immense parallel-processing capabilities, have enabled the training of large-scale AI models that were previously unimaginable. Rapid adoption and integration of these new GPUs by companies like Microsoft and Google since the beginning of 2023 has only further strengthened Nvidia's lead in parallel computing.

Bigger Than AI: How Parallel Computing Will Change the World

Nvidia's advancements in parallel computing extend far beyond AI. It might sound like an exaggeration, but we believe these advancements could be as transformative to the global economy as the printing press, the Industrial Revolution, or the rise of the internet.

Consider Gutenberg's printing press, which revolutionized the spread of knowledge by making books widely accessible. Nvidia's progress in parallel computing has similarly democratized access to high-performance computing capabilities, giving researchers, scientists, and entrepreneurs – as well as normal people – access to computational power that was previously reserved for large supercomputing facilities.

The Industrial Revolution overhauled manufacturing through mechanization, boosting productivity and economic growth. Parallel computing could usher in similar improvements by enabling intelligent systems, autonomous vehicles, advanced simulations, and smart robots that can work 24/7. This could drive unimaginable productivity and efficiency gains across the economy in the coming decades.

And the internet transformed communication, connectivity, and information sharing, leading to profound social and economic changes. Parallel-computing advancements could ultimately enable new forms of human-to-human and human-to-machine communication and connectivity – such as the metaverse or direct brain-computer interfaces – to dramatically add new possibilities to business, to society, and to human interaction.

Finding Ways to Profit From the Parallel Computing Revolution

These advancements in computing could ultimately create trillions of dollars of wealth in the decades ahead. However, profiting from this great leap – via direct investment in Nvidia stock or in a handful of other important ancillary companies – will require discipline and patience.

At its current $135 share price, Nvidia commands a market capitalization of $3 trillion. This makes it the largest company in the global economy. Despite its gargantuan valuation, Nvidia shares are not outrageously priced, given two key assumptions: 1) demand for its GPUs can continue producing robust revenue growth exceeding 30% annually, and 2) Nvidia can maintain its world-class 55% profit margins.

If these assumptions hold true, Nvidia should generate roughly $160 billion in revenue in 2025 and $35 in earnings per share – roughly double what it generated in the last 12 months on both metrics. That puts its forward price-to-earnings ratio at just 34x, compared with just over 20x for the S&P 500. Given its dominant market position, growth, and profitability, this is arguably not an extreme valuation, considering it's one of the best businesses on earth.

That said, investors should be wary of the risks embedded in the two assumptions laid out above. On the growth assumption, it's easy to imagine a scenario where Nvidia's biggest customers – Microsoft (MSFT), Meta Platforms (META), and Amazon (AMZN) – suffer from a slowdown in their businesses, should the U.S. or the global economy enter a recession.

Demand for cloud computing and digital advertising, two tech sectors driving economic growth at the moment, will not be immune from a broader slowdown. If that occurs, these companies will likely pull back on their capital spending – which means less money flowing to Nvidia for its GPUs. It's also worth noting these same companies each have their own development programs in the works to produce their own GPUs to compete against Nvidia.

In addition, it's important to note that Nvidia doesn't actually manufacture its own chips – it designs them and relies on others to actually make them. So it's conceivable that Nvidia's key suppliers – companies like Taiwan Semiconductor Manufacturing (which we'll cover in next week's issue) – could begin charging Nvidia higher prices to manufacture its chips.In June news broke that TSM was in talks with Nvidia to do exactly that. Given that TSM controls 90% of the global manufacturing capacity for high-end GPU manufacturing, the company holds a powerful negotiating hand, and could begin chipping away at Nvidia's margins.

Where the Parallel Computing Revolution Is Headed Next

This current scenario reminds us of the work of the great author, technology advocate, and free-market thinker George Gilder – an early prophet of the internet who correctly predicted the rise of many of today's most successful technology companies.

The problem with visionaries is that they can see too far too fast – and the market cannot always keep up. Virtually all of the technology stocks Gilder recommended in the late 1990s produced market-beating returns over several decades – but only after first shooting up like flares before falling back to Earth. Investors who bought in when Gilder initially recommended these companies first suffered through gut-wrenching drawdowns of 70% to 90%.

Internet leader Microsoft (MSFT) is a quintessential example. Microsoft's share price fell more than 65% as the dot-com boom turned to bust in 2000. Investors who owned MSFT at that time would have to wait nearly 17 years to see the stock return to those prior highs, even as the company's revenues continued growing by double-digit rates each year.

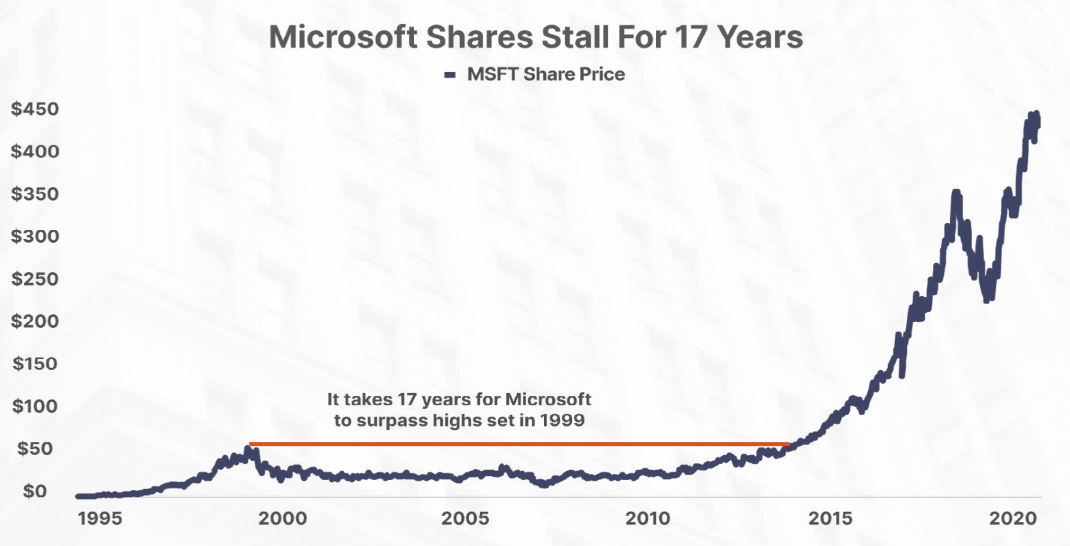

Source: Bloomberg

In reality, the best time to invest in MSFT was not in the late 1990s but rather in late 2000, after shares had plunged and most investors had given up.

We believe a similar dynamic is playing out with many of the AI and machine-learning companies utilizing the chips and the parallel-computing trend started by Nvidia. The market is clearly in a bubble phase. However, there is simply no way to know how long it will continue – or what the pin that finally pops it will be.

To return to the dot-com bubble analogy, 2024 could be equivalent to 1996, which kicked off several years of double-digit gains before the broad market reached its ultimate peak; or 1997-1998, when volatility exploded higher following the Asian Financial Crisis, yet the biggest gains were still to come; or 1999-2000, when many internet stocks were already peaking, and a prolonged bear market was just around the corner.

PART 2: NVIDIA'S SILENT PARTNERS

Nvidia is the epicenter of the Parallel Processing Revolution… but it can't function alone. The parallel-processing giant relies on a supply chain of components made by lesser-known companies… each one forming a vital building block of the industry.

In the first of the reports in this section, **The New OPEC**, you'll get insight into two of these partners… companies that control critical chip production resources and make what Nvidia does possible.

The first company, **ASML (Nasdaq: ASML)**, is the global leader in photolithography machines used in semiconductor manufacturing. With a market cap of $424 billion, it generated $28 billion in revenue in 2023, a 29% increase over 2022.

This company's lithography machines are the only devices in the world capable of manufacturing the special computer chips that power Nvidia GPUs, along with every other high-end computer chip on the market. Today, ASML owns 90% of the global market share for all lithography machines used in semiconductor manufacturing.

The next building block, **Taiwan Semiconductor Manufacturing (NYSE: TSM)**, ihas a monopoly on semiconductor chip manufacturing. This hugely successful Taiwanese company controls over 90% of the supply of advanced semiconductors worldwide. It has long-standing partnerships to supply advanced chips for virtually all of the world's largest semiconductor companies. Its key customers include Apple, Nvidia, Qualcomm, Broadcom, and Advanced Micro Devices ("AMD").

That means this company represents a bigger geopolitical choke point – on a product that's arguably more important to the global economy than oil – than all of OPEC, which controls just 40% of the world's oil supply.

The second report in this section, **The Missing Link**, explores an essential piece of the parallel-processing supply chain…

Arm Holdings PLC–ADR (Nasdaq: ARM) creates the blueprints for semiconductor chips. Arm is the only company in the world that owns the intellectual property underlying the energy-efficient Grace CPU, which is a crucial component for Nvidia. It licenses these blueprints to other semiconductor companies that design chips based on its proprietary technology. And it collects a perpetual revenue stream for as long as the chip remains in production, which can be for decades.

This makes Arm one of the most capital efficient businesses in the world, with the company spending less than $100 million in capex each year to produce over $3 billion in revenue. As a result, the company generates roughly $1 billion in annual free cash flow, for a stellar 31% margin.

Porter & Co.

THE BIG SECRET ON WALL STREET

The New OPEC

- How Two Companies Control the World's Most Valuable Resource
- The Ultimate "Forever Stocks" for Today's Computing Revolution

FROM THE DESK OF PORTER STANSBERRY

The New OPEC

How Two Companies Control the World's Most Valuable Resource

The Ultimate "Forever Stocks" for Today's Computing Revolution

Crash!

On a micro scale, it could have been the sound of nuclear fission at the heart of an atomic weapon.

On the streets of Albuquerque one December day in 1945, it was the sound of two inept young spies bumping into each other.

Nineteen-year-old Ted Hall and 20-year-old Saville Sax, two young Communist-leaning Americans, had accepted a dangerous assignment from the Soviet Union… steal the plans for America's top-secret nuclear bomb, code-named the Manhattan Project.

After four years of development at classified locations in the U.S. and England, the bomb was nearly complete, ready to deliver the final blow to the faltering Axis powers at the end of World War II.

And Soviet dictator Joseph Stalin felt a little left out.

Though Russia was allied with America and Britain, the two Western nations didn't trust the Union of Soviet Socialist Republics *quite* enough to share their atomic secrets. But "Uncle Joe" Stalin wanted a nuclear bomb for Mother Russia… and what Stalin wanted, he usually got.

So as American and British scientists fine-tuned the fission and fusion of the atom, the Soviets infiltrated the Manhattan Project labs with a ring of carefully selected informants… known today as the "atomic spies."

Ted, a Harvard graduate by 18, and the youngest physicist working on the bomb, was an obvious target for the Soviets. He'd put on what he called "pink-colored glasses" during his time at college, and believed nuclear weaponry should be shared equally with all.

Ted had something *extra* valuable to share with Russia: he'd personally designed the central combustion mechanism of the "Fat Man" plutonium bomb that would soon destroy Nagasaki, Japan.

As the informant, Ted (code-named "Young") passed critical intel to Saville, the courier (code-named "Old"), who then headed back to the spies' Russian handler in New York.

Ted was young indeed. And his idea of espionage came straight out of the funny papers.

No one saw Ted approach the Albuquerque meetup from the wrong direction and smash into his contact… or the cartoon-style handoff where Saville stuffed the blueprint for “Fat Man” into his shoe.

Though the handoff wasn’t a masterclass in spycraft, the traitorous teen still single-handedly gifted the Russians their own nuclear weapon. In 1949, the Soviets tested their first homegrown plutonium bomb… using the same groundbreaking design that Ted had slipped to Saville.

Ultimately, Ted also gave the world the gift of the Cold War, where America and Russia spied on each other obsessively and amassed ever-growing piles of deadly nuclear technology during the protracted middle part of the last century.

In an interview with *The New York Times* in 1997, a 71-year-old Ted Hall suggested that he’d changed the course of history. And he was, by and large, proud that he’d stopped American nuclear monopoly in its tracks. “Maybe the ‘course of history,’ if unchanged, could have led to atomic war in the past 50 years – for example the bomb might have been dropped on China in 1949 or the early 50’s,” he said. “Well, if I helped to prevent that, I accept the charge.”

Where there are secrets, there are always spies. Carefully guarded inventions invariably spark competition and chicanery. And whenever one group monopolizes a game-changing technology, some kind of arms race is sure to follow.

Human civilization is built of space races, arms races, oil races – a never-ending cycle of cold and hot wars as superpowers fight for limited resources. In the early decades of the 21st century, in addition to our perennial global tug of war over weapons and energy, we’ve added a new conflict into the mix: the chip race.

Complete with its own tight monopolies and industrial spies, this new arms race centers on the global battle to control advanced semiconductors, known as sub-10 nanometer (“nm”) chips. These tiny slivers of metal – and the factories that build them – are the cornerstone of the global economy, powering everything from smartphones to data centers to missile guidance systems, not to mention the accelerating artificial-intelligence (“AI”) industry.

Remarkably (and similar to the American-British chokehold on nuclear technology during the 1940s), the supply chain for the massive semiconductor industry – which generates about $600 billion in revenue per year – rests in the hands of just two companies… making the rest of the world almost totally dependent on them.

Neither of these companies are in China. And oh boy, do the Chinese want a piece of them…

Last year, China spent $390 billion importing semiconductor chips and manufacturing equipment – 15% more than what it spent importing oil. Right now, China manufactures a few chips of its own, but none of the most advanced variety. So it's scrambling madly to beg, borrow, and steal intellectual property ("IP") from chipmakers around the world... particularly, the two companies with the virtual chip monopoly.

Not surprisingly, China currently has its own semiconductor version of atomic spies – and at least one of them got off scot-free.

Ted was young indeed. And his idea of espionage came straight out of the funny papers.

No one saw Ted approach the Albuquerque meetup from the wrong direction and smash into his contact... or the cartoon-style handoff where Saville stuffed the blueprint for "Fat Man" into his shoe.

Though the handoff wasn't a masterclass in spycraft, the traitorous teen still single-handedly gifted the Russians their own nuclear weapon. In 1949, the Soviets tested their first homegrown plutonium bomb... using the same groundbreaking design that Ted had slipped to Saville.

Ultimately, Ted also gave the world the gift of the Cold War, where America and Russia spied on each other obsessively and amassed ever-growing piles of deadly nuclear technology during the protracted middle part of the last century.

In an interview with *The New York Times* in 1997, a 71-year-old Ted Hall suggested that he'd changed the course of history. And he was, by and large, proud that he'd stopped American nuclear monopoly in its tracks. "Maybe the 'course of history,' if unchanged, could have led to atomic war in the past 50 years – for example the bomb might have been dropped on China in 1949 or the early 50's," he said. "Well, if I helped to prevent that, I accept the charge."

Where there are secrets, there are always spies. Carefully guarded inventions invariably spark competition and chicanery. And whenever one group monopolizes a game-changing technology, some kind of arms race is sure to follow.

Human civilization is built of space races, arms races, oil races – a never-ending cycle of cold and hot wars as superpowers fight for limited resources. In the early decades of the 21st century, in addition to our perennial global tug of war over weapons and energy, we've added a new conflict into the mix: the chip race.

Complete with its own tight monopolies and industrial spies, this new arms race centers on the global battle to control advanced semiconductors, known as sub-10 nanometer ("nm") chips. These tiny slivers of metal – and the factories that build them – are the cornerstone of the global economy, powering everything from smartphones to data centers to missile guidance systems, not to mention the accelerating artificial-intelligence ("AI") industry.

Remarkably (and similar to the American-British chokehold on nuclear technology during the 1940s), the supply chain for the massive semiconductor industry –

which generates about $600 billion in revenue per year – rests in the hands of just two companies… making the rest of the world almost totally dependent on them.

Neither of these companies are in China. And oh boy, do the Chinese want a piece of them…

In 2023, China spent $390 billion importing semiconductor chips and manufacturing equipment – 15% more than what it spent importing oil. Right now, China manufactures a few chips of its own, but none of the most advanced variety. So it's scrambling madly to beg, borrow, and steal intellectual property ("IP") from chipmakers around the world… particularly, the two companies with the virtual chip monopoly.

Not surprisingly, China currently has its own semiconductor version of atomic spies – and at least one of them got off scot-free.

Mutually Assured Dinero

Of the hundreds of Communist sympathizers in the U.S. who were Soviet informants during the Manhattan Project, only a few, very high-profile spies were ever tried and convicted. Klaus Fuchs, a high-level turncoat scientist, got nine years in jail, and Julius and Ethel Rosenberg notoriously received the death penalty for "conspiracy to commit espionage" in 1953.

Declassified Russian documents show that countless code-named others slipped away: "Fogel," "Mar," "Quantum," "Eric." Melita Norwood, an unassuming secretary who filched bomb blueprints throughout her 40-year career and only fessed up after her retirement. And, of course, Ted Hall ("Young") – whose real name (presumably accidentally) appeared once in his Soviet handlers' files.

Ted was questioned by U.S. government authorities a few times during the war, but was never apprehended – and he went on to get married and live an unremarkable life as a researcher in biophysics, granting a few cagy interviews toward the end of his life. The most he ever outright admitted to was being "worried about the dangers of an American monopoly of atomic weapons."

It seems that Chinese engineer Zongchang Yu has managed so far to wriggle through the net in much the same way.

Yu weaseled his way into a top spot at one of the two companies that controls the chip industry – and then left the company in 2012, taking a sizable chunk of its proprietary software with him. Two years later, Yu launched Dongfang, a now-successful Beijing-based semiconductor company.

There's still a warrant out for Yu's arrest in his former adopted state of California on a count of intellectual property theft… but the Santa Clara County sheriff is unlikely to find him behind the Bamboo Curtain. No doubt, like Ted, Yu is "worried about a monopoly" on semiconductor chips.

To be fair, China is not the only country trying to snatch a helping of chips. Neither of the monopolizing companies is American – so, in order to keep up, America is pouring billions of dollars into developing its own semiconductor industry.

In August 2022, Congress passed and President Joe Biden signed the cleverly-named CHIPS Act (Creating Helpful Incentives to Produce Semiconductors). This act – the largest-ever domestic infrastructure bill aimed at a single industry – authorizes roughly $280 billion to fund research and manufacturing capacity for domestic semiconductor production in the U.S. This government funding, along with tens of billions from the private sector, aims to reverse a four-decade long trend of U.S. chipmakers outsourcing production.

However, as we'll show in this issue, despite the flood of Chinese and American investment dollars and blatant Chinese IP theft, neither country has succeeded in displacing the two choke points in the global chip-production supply chain.

The reason: these companies control the most sophisticated manufacturing operations in human history. They maintain such a powerful, entrenched advantage over the rest of the industry that the best efforts of the two largest governments in the world have failed to replicate them.

And all evidence suggests that these two firms will continue to dominate the global semiconductor supply chain for years to come. These companies – which, although they are not headquartered in the U.S., are publicly traded on the Nasdaq and NYSE – are the two best ways for investors to capitalize on today's parallel-processing revolution.

To bastardize another arms race term, it's mutually assured dinero.

Let's get started with the largest monopoly of them all – the company with 100% market-share dominance in the critical machines used to make today's most advanced semiconductors.

Leaving Moore's Law in the Dust

Founded in the Netherlands in 1984, **ASML (Nasdaq: ASML)** is the global leader in photolithography machines used in semiconductor manufacturing. It generated $28 billion in revenue in 2023, a 29% increase over 2022.

Photolithography machines use light energy to etch intricate patterns onto the silicon wafers of semiconductors. These patterns enable the precise placement of components onto computer chips – including transistors, the key building block of electronic devices.

First, though, a bit of background on transistors and their critical role in driving advancements in computing power...

Transistors exist in one of two states: on or off, depending on the flow of electricity through a computer chip. This on/off state translates into a 0 or 1, which forms the basis of the binary code instructions in all computer-programming languages.

The more transistors on a computer chip, the more binary code it can process, and thus the greater its computing power.

As we previously wrote, the co-founder of U.S. chipmaker Intel, Gordon Moore, noticed in the 1970s that the number of transistors going into cutting-edge microprocessors was doubling every 18 to 24 months – and so the processing speed was increasing at similar rates. The observation became known as Moore's Law, and the advancements in computer processing speeds has largely followed this projection over the last five decades. This exponential increase in computing power has made modern-day life possible.

While chipmakers get much of the credit for engineering and manufacturing ever-more advanced semiconductors, none of this progress would exist without the advancements in photolithography equipment. By finding new ways to etch smaller, more precise patterns into silicon wafers, lithographic innovation has enabled manufacturers to continue packing more transistors into semiconductors.

When the first silicon-based semiconductor was introduced in the 1960s, it contained four transistors measuring 40 micrometers (one micrometer is one-millionth of a meter), or about half the width of a human hair. By the 1990s, transistor sizes had fallen to hundreds of nanometers (one nanometer is one-billionth of a meter), allowing chipmakers to pack over 10 million transistors into a single semiconductor. Since then, transistor sizes have fallen by more than 100-fold to less than 10 nanometers today – or about the width of a cell membrane.

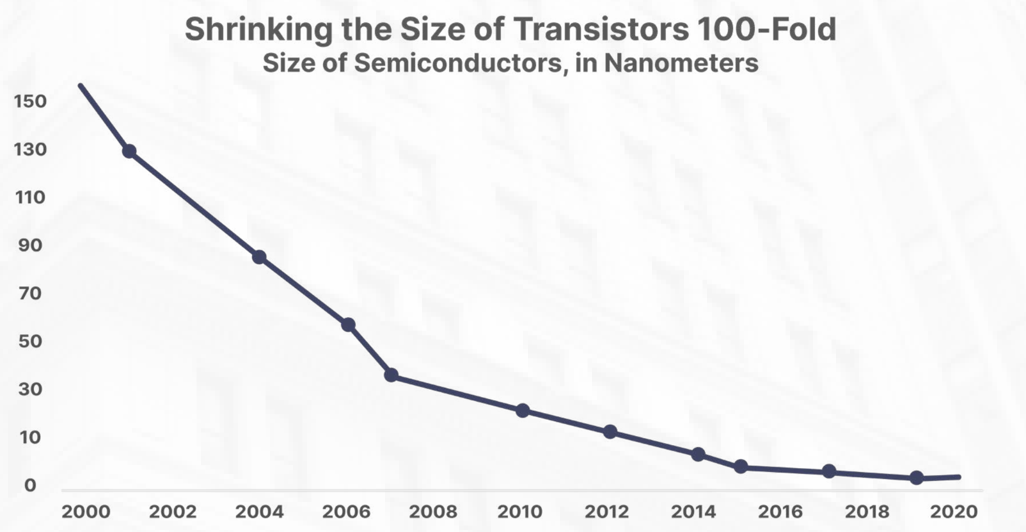

The latest class of sub-10 nm semiconductors powers today's parallel-processing revolution. It's how Nvidia designed its latest Blackwell B100 GPU with an incredible 208 billion transistors packed into a single chip – the most of any GPU chip in the world. A GPU, as we discussed in detail last week, is a graphics processing unit – a specialized computer processor originally designed to speed up graphics rendering in video games.

ASML's lithography machines are the only devices in the world capable of manufacturing the sub-10 nm chips powering Nvidia GPUs, along with every other high-end computer chip on the market. The story of how it got here provides the key context for understanding why ASML will remain a monopoly for years to come.

Early Success Stalled by a Technological Wall

When ASML first entered the lithography market in 1985, it was a small fish in a big pond. Japanese firms like Canon and Nikon dominated the industry, and ASML struggled to compete against these larger and more established rivals.

ASML's first breakthrough came in 1991, with the release of its cutting-edge PAS 5500 lithography machine. ASML struck a long-term deal to supply the PAS 5500 for Micron Technology, at the time one of the world's largest memory and storage chipmakers. The deal turned ASML from a struggling upstart into a respected name in the industry. The deal also transformed ASML's financials from showing consistent net losses to generating over $100 million in net income every year by the mid-1990s. The company reinvested heavily into research and development (R&D) to launch a series of newer, more advanced machines throughout the decade.

Lithography machine makers like ASML have two primary modes of innovation for etching ever-smaller transistor sizes onto silicon wafers.

The first comes from tapping into new sources of lower-wavelength light energy. Think of shorter wavelength light as providing a thinner paint brush to imprint finer details onto a canvas – or silicon wafers in the case of chip making. When the first semiconductors were created in the 1960s, lithography devices used visible light in the 435 nm spectrum. From there, the industry progressed toward ultraviolet light, in the 365 nm spectrum, and then deep ultraviolet systems, at 193 nm, by the 1990s.

The second key innovation for shrinking light wavelengths comes from advancements in optical systems. Lithography machine makers use complex optics systems containing a series of lenses to focus and reduce the source light down to even-thinner wavelengths.

Throughout the 1990s, ASML rose to the forefront of the industry by implementing cutting-edge optics systems, first tapping into ultraviolet light in the 365 nm range then using advanced lasers for tapping into the 193 nm range.

However, by the late 1990s, ASML and other lithography leaders ran into a technological wall. The industry had exhausted its ability to harness shorter wavelengths of light energy, and advancements in optics could only go so far in shaving deep ultraviolet light down to fewer nanometers of functional wavelengths.

The next step involved crossing a major chasm into extreme ultraviolet light ("EUV"). This promised to shrink the silicon-board paint brush by 90% from the 193 nm wavelengths in the deep ultraviolet spectrum down to 13.5 nm of EUV light. But tapping into this new frontier presented massive technical challenges that defeated virtually every company that attempted it.

Crossing the EUV Chasm With Help From Uncle Sam

The first major obstacle to unlocking EUV lithography technology was producing a 13.5 nm source light. At the time, no commercial laser system existed for reliably producing EUV light at the demanding levels of precision required in lithography machines.

Another major complication was the fact that EUV light is absorbed by glass and air. This meant the traditional industry approach of using glass lenses to focus the light beams was not workable. Instead, engineers would need to use mirrors to direct the EUV light beam onto silicon wafers. But the 13.5 nm wavelength of EUV light is so small that the existing mirrors available wouldn't reflect it. Afinal challenge was that the whole thing needed to operate within the confines of a sealed vacuum chamber, to eliminate interaction with ambient air (which is air in its natural state, which would absorb the EUV energy).

The scientific and engineering hurdles were so daunting that two of the leading lithography companies of the day, Canon and Nikon, gave up on EUV technology in the 1990s. By this point, many considered the development of EUV lithography to be impossible. Some even predicted that the stalled progress in lithography advancement would mean the end of Moore's Law.

In 1997, the U.S. government stepped in to help the industry solve the daunting technological challenges of EUV lithography. The Department of Energy ("DoE") funded a six-year EUV research program, managed through a public-private partnership between top university laboratories and industry leaders. The U.S. government retained control over any technology developed through the partnership, which it owned though an entity called the Extreme Ultraviolet Limited Liability Company (EUV LLC). This meant that any company that wanted to use the technology developed by EUV LLC would first need to get approval from the U.S. Congress.

The EUV LLC initiative achieved numerous technological milestones that would later lead to the development of the first commercial EUV lithography machines. But the U.S. government was very selective in granting licenses to this critical technology. Specifically, the DoE denied access to Nikon and Canon.

One of the companies that participated in the EUV LLC initiative, and thus gained access to the technology license, was Silicon Valley Group ("SVG"). SVG was a U.S.

lithography equipment maker based in San Jose, California, that ASML acquired in 2000 (and for a strategic reason we will explain later), giving ASML ownership of the technology license. Over the next six years, the company invested heavily into advancing the learnings from EUV LLC. Then in 2006, ASML unveiled the world's first working prototype of an EUV photolithography machine.

But producing a successful prototype was only the beginning of a long slog toward commercializing the technology. It took ASML another seven years – and billions more in R&D – to produce the world's first commercially available EUV lithography machines. Fast forward more than a decade, and ASML is the only company in the world that has developed usable EUV lithography.

To understand why this technology required such a painstakingly long development period, let's explore what have been called "the most sophisticated machines ever manufactured."

Small Transistors, Giant Machines

ASML's latest product iteration, released in 2023, is called a High-NA-EUV lithography machine. The High-NA stands for high numerical aperture, which describes the cutting-edge optics system used to shrink EUV light down to an 8 nm wavelength. This reflects a 1.7x improvement from ASML's original 13.5 nm EUV lithography machines, first released in 2013.

Creating these nano-sized light wavelengths requires a massive machine containing over 100,000 parts and weighing 165 tons – the equivalent heft of two commercial jet airplanes. When ASML shipped its first High-NA-EUV machine to an Intel factory in Oregon last December, the parts alone filled the cargo holds of four 747s. And once the parts arrived, it took 250 engineers six months to install the mammoth machine (pictured below).machine (pictured below).

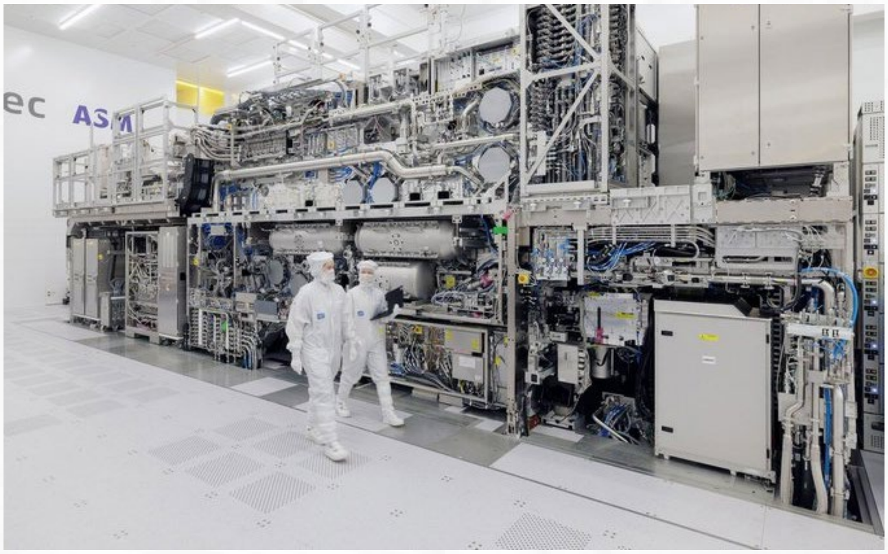

Properly describing the intricacy of these machines would require a full-length dissertation.

For a brief bit of insight, consider what goes into the production of the EUV source light alone. ASML, working with its key suppliers, had to create one of the most sophisticated laser systems in the world to generate this light source.

The process begins with a device that shoots tiny droplets of molten tin (25 microns in diameter, or about a third of the width of a human hair) through the air at 230 feet per second. As the drops fall, a laser beam strikes the tin molecule at precisely the right angle to flatten it into the intended shape. Next, a second laser beam vaporizes the tin into a superheated plasma, reaching a temperature of 220,000 degrees Celsius (40 times hotter than the sun) to generate the EUV light. The process is repeated 50,000 times each second, requiring unimaginable levels of speed and precision. The calculations used to zap the tin molecules in mid-air involve more precision than those used to navigate the Apollo lunar module to the moon.

Once the EUV light is created, it must be directed toward precise, nano-sized targets on the silicon wafer through a complex mirror system. The EUV light wavelength is so small that typical high-tech mirrors won't reflect it. So ASML had to work with its key optics supplier – German company Zeiss – to create a custom-designed mirror just for ASML machines. The manufacturing process uses a high-precision ion beam that shaves the mirror surface one individual molecule at a time, making it what ASML calls "the most precise mirror in the world."

The diagram below shows the complex gauntlet of mirrors the machine uses to direct EUV light onto a silicon board, etching out patterns within a quarter of a nanometer's precision (about the width of an individual atom):

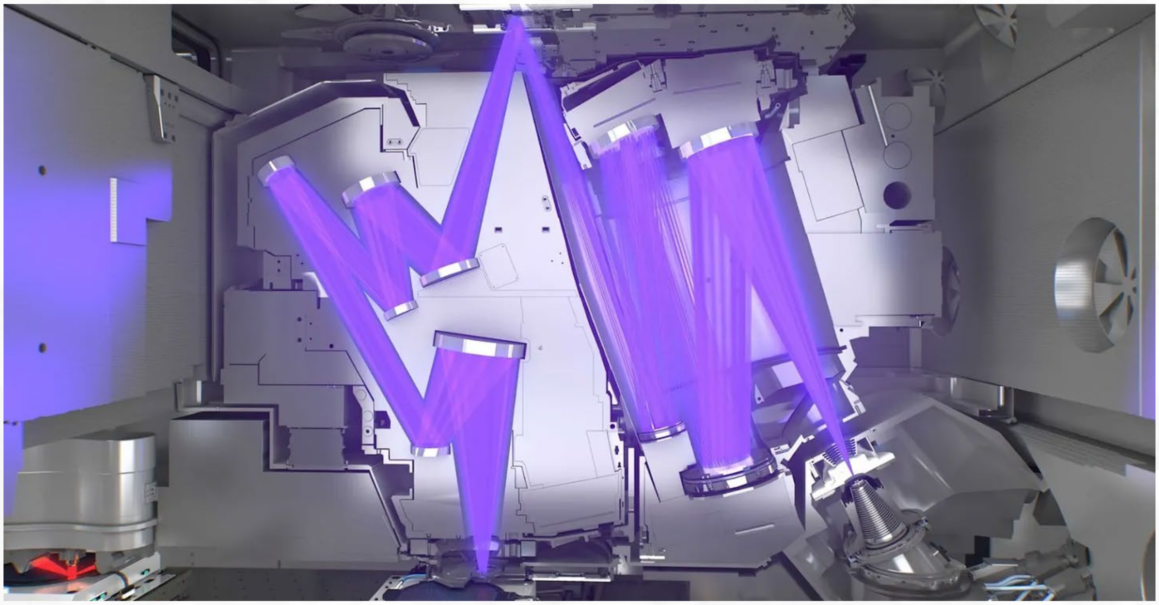

This is just one aspect of the complexity that goes into ASML's machines that etch the intricate patterns in today's most advanced 2 nm to 3 nm semiconductor chips.

Overruling Moore's Law

While companies like **Nvidia (NVDA)** deserve much of the credit for engineering today's cutting-edge sub-10 nm chips, they would be impossible to manufacture without ASML.

Leading tech publication *MIT Technology Review* has heralded ASML's EUV lithography equipment as "the machine that saved Moore's Law."

To be clear, ASML's ongoing advancements in EUV lithography saved the idea of Moore's Law and its notion of measuring constant improvements in processing speed. But the original Moore's Law computation has been left in the dust.

Consider the trajectory of computing power in Nvidia's cutting edge GPU chips in recent years. The standard measure of computational capacity is a term known as "floating point operations per second" ("FLOPS"), which refers to the number of arithmetic calculations a chip can perform per second. The power of today's most advanced chips are measured in Teraflops ("TF"), where one TF equals one trillion FLOPS.

In just the last eight years, the computational capacity of Nvidia's GPU chips, all of which are manufactured by ASML machines, have increased 1000-fold from 19 TFs to 20,000 TFs. That's four times the 256x rate of increase suggested by Moore's Law over the same period:

Source: analyticsindiamag.com

Today, ASML controls 90% of the global market share for all lithography machines used in semiconductor manufacturing. But most importantly, it sells 100% of the EUV lithography machines used to make sub-10 nm semiconductors. The demand for these chips goes beyond Nvidia GPUs and the AI arms race. These semiconductors power all of the latest generation of electronics used in consumer devices (smartphones, tablets, and PCs), industrial applications (data centers and 5G telecommunication networks), and in the military (drones, fighter jets, and missile guidance systems).

So while ASML is one of the biggest beneficiaries from the rise of parallel computing, its revenue stream is far more resilient than most other companies riding this boom. Given its monopoly position as the sole lithography supplier to high-end chipmakers across every end market, ASML enjoys a steady source of new demand from the rising tide of semiconductor manufacturing across the board.

That's how the company has increased its revenue over the last decade at a compounded annual growth rate ("CAGR") of 17%, regardless of the booms and busts in different chip segments (i.e., crypto mining, metaverse, and now AI and parallel computing):

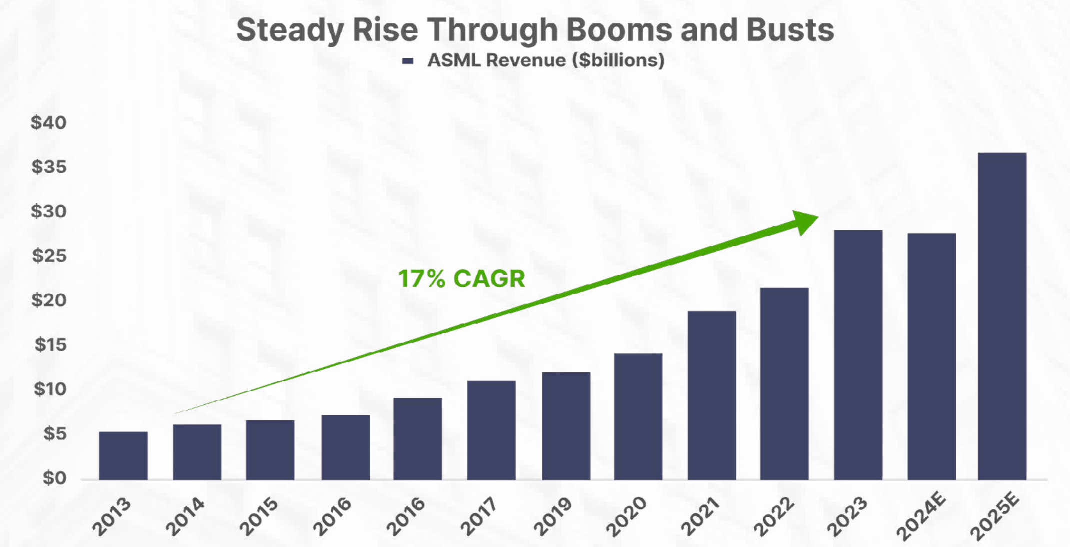

Source: Bloomberg

There's one key factor that instills great confidence in ASML's future growth prospects, regardless of which companies win the AI arms race, or how the trends in semiconductor technology evolves. That one key factor is its virtually impenetrable competitive moat that continues growing wider.

A Dominant and Growing Competitive Moat

When ASML first entered the lithography market, it was playing from a position of weakness against the industry giants Nikon and Canon. Even after ASML introduced a line of successful products in the mid-1990s, the two companies together generated 10 times the earnings of ASML. But since commercializing EUV lithography, ASML has produced a growing stream of profits that it has continuously recycled into R&D. Today, as the world's dominant supplier of lithography machines, ASML generates four times as much income as those top two rivals combined:

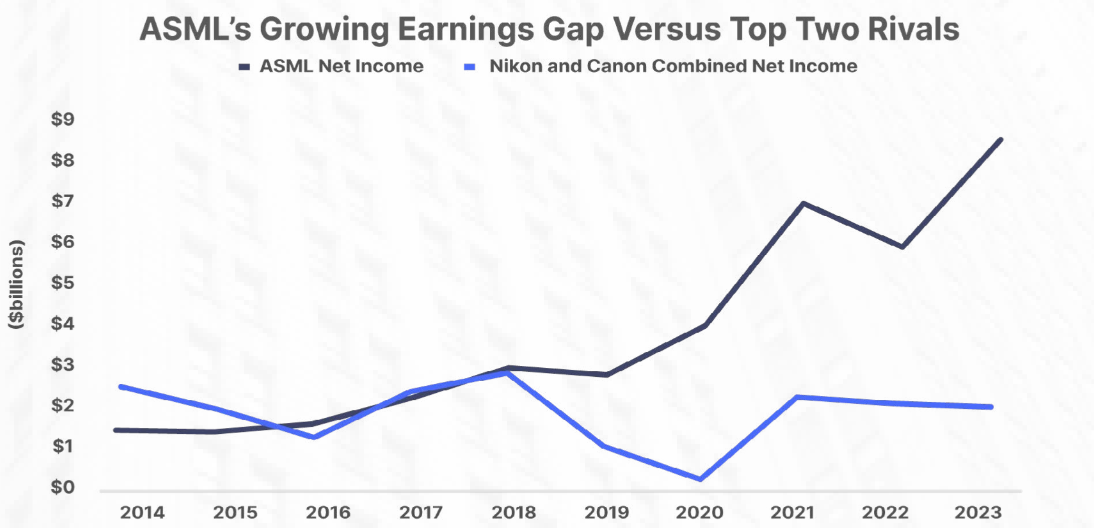

Source: Bloomberg

ASML is now playing from a position of incredible strength. Its unmatched earnings power gives it the ability to pour $1 billion a quarter into R&D, or double what its next closest competitor, Nikon, invests. That's how the company continues introducing cutting-edge EUV lithography machines, while the rest of the industry struggles to even gain a foothold into the market.

But ASML's R&D spending is only the first layer of its competitive moat. Even when companies with bigger R&D budgets have attempted to enter the market, they've failed.

This most notable example is Intel – one of the world's largest chipmakers – that tried mastering EUV lithography for its own semiconductor manufacturing lines in 2012. At the time, Intel's $5 billion R&D budget dwarfed the $600 million ASML invested each year into R&D. In 2012, Intel announced a $4 billion commitment to developing its own EUV lithography machines.

Intel was part of the original consortium of private companies that partnered with the U.S. Department of Energy as part of the EUV LLC initiative in the 1990s. This meant Intel had access to the same technological starting point as ASML, and a significantly larger budget.

And yet, Intel gave up on its EUV program. Today, Intel is now a customer of ASML – paying ASML $380 million for just one High-NA-EUV machine in 2023.

Even the Chinese company Dongfang, created in 2014 by Zongchang Yu – the former ASML engineer who orchestrated the theft of all 2 million lines of source code used in ASML's EUV lithography machines – has failed to produce a commercially viable EUV system. And that's despite the backing of the Chinese government.

Thus, the question at hand: why have no other companies successfully replicated ASML's EUV technology – even well-financed competitors working with bigger R&D budgets and similar technological starting points and those using an exact copy of ASML's stolen intellectual property?

ASML's Big Secret: A Dominant Supply Chain

ASML has done far more than master an incredibly hard-to-implement technology by itself. The company has amassed a complex supplier network spanning over 5,000 companies that work hand-in-hand with ASML to produce the more than 100,000 components in each EUV machine. Many of these suppliers are highly-specialized equipment makers that are themselves monopoly providers of equipment. In some cases, ASML has exclusive partnerships that prohibit doing business with the competition.

Consider the case of optics maker Zeiss. Zeiss, as mentioned earlier, worked with ASML to develop the custom mirrors capable of reflecting EUV light. ASML has maintained a long-running partnership with Zeiss. The relationship frayed in the late 1990s when ASML discovered that Zeiss was selling optics equipment to competitor SVG. ASML responded by acquiring SVG and locking in an exclusive relationship with Zeiss. In 2016, ASML took a 25% equity stake in Zeiss, and as part of the deal, gave Zeiss CEO Peter Grassmann a seat on ASML's board. Today, ASML refers to the relationship as "two companies, one business."

When ASML can't control suppliers, or when quality standards disappoint, it simply buys them. Author Chris Miller describes these aggressive tactics in his best-selling book Chip War: The Fight for the World's Critical Technology. Miller explains that Peter Wennink, who served as ASML's CEO from 2013 - 2024, has been known to warn suppliers: "If you don't behave, we're going to buy you."

One high-profile purchase by ASML was of EUV-light-source manufacturer Cymer, for $2.5 billion in 2012. That was the same year Intel announced its $4 billion commitment to develop its own EUV lithography system. Cymer's laser technology was considered so critical to EUV development that an industry analyst with investment bank RBC Capital Markets noted at the time: "This is a one-of-a-kind deal... EUV is pretty much down to ASML."

And therein lies the big secret to ASML's iron grip on EUV lithography. With more than 100,000 parts needed to manufacture these systems, no one company can do it all on its own. ASML has a stronghold on the key choke points of the supply chain needed to produce these machines. This includes a tight strategic relationship with Zeiss, the only optics company capable of making the mirrors that reflect EUV wavelength light. And through its acquisition of Cymer in 2012, ASML owns the only company that can produce critical components for EUV light-source generation.

These are just two examples that illustrate the key point: replicating ASML's success in EUV lithography machines requires much more than a giant R&D budget, or simply stealing the intellectual property.

An analogy would be a time traveler going back to the 1700s and dropping off the blueprints for a Ford F-150 pickup. Even if a person from that era could decipher the diagrams, they'd have no way of securing the machined metal and other parts needed to build the vehicle. Without a vast network of component suppliers, the blueprints would be worthless.

The same is true for ASML's competitors. Any new entrant into this market would need to painstakingly replicate ASML's 5,000-plus network of key suppliers. But ASML has already taken many of the critical players in this supply chain off the market, either through acquisitions or exclusive partnerships.

There's one final aspect of ASML's business that's hard to replicate: the ongoing service and support required to maintain the company's complex machines. Each lithography machine that ASML ships comes with a full-time support staff of technical experts that remain onsite for the full length of the machine's life (typically around 20 years). These technicians do everything from routine maintenance to replacing broken parts, as well as continuously fine-tuning the precision machine for optimal performance.

Training a workforce with the high-tech skills needed to provide this support only comes with years of hands-on experience working directly with these machines - a process that would take years to develop, even after a competitor successfully mastered EUV lithography technology.

Another advantage of this service – it provides ASML with a steady source of recurring revenue, which has historically made up between 20% to 30% of its business. This creates a high degree of stability in ASML's business, by providing a source of ongoing returns from each machine for up to 20 years after the original sale. And this high level of service is critical for ASML's customers, by ensuring they maximize the performance and run-time out of each machine they buy. Any new competitor hoping to enter the market would have to replicate this critical service feature of the business.

When viewed through this lens, it's easy to see how ASML has established a monopoly position in EUV lithography machines. Through its acquisition of SVG in 2000, it secured an early advantage over its two key rivals, Nikon and Canon, by tapping into the exclusive technology developed by the EUV LLC public-private partnership. From there, it poured billions into advancing the technology and developing the first working prototype in 2006.

By the time other competitive threats emerged, including Intel in 2012, ASML had secured major choke points of the supply chain required to produce these incredibly complex machines. This has made it virtually impossible for its competitors to enter the market, allowing ASML to further extend its advantages in technology, supply-chain dominance, and service. From this leading position,it's very difficult to imagine a scenario where a competitor displaces ASML.

The Ultimate Forever Stock

ASML's dominant competitive position as the sole supplier of EUV lithography machines makes it the ultimate "forever stock." Its competitive position has been secured through 30 years of R&D investment and the cultivation of its intricate supplier network.

The company's products will remain in high demand so long as the world continues requiring higher volumes of advanced sub-10 nm semiconductors. This is a trend that will likely continue indefinitely, regardless of which companies and technologies rise and fall along the way.

ASML's monopoly position ensures a long runway of growth, and also gives it incredible pricing power. Nowhere is this more evident than in the hefty $380 million price tag of its latest High-NA-EUV machines.

At a May technology symposium in Amsterdam, a senior executive at Taiwan Semiconductor ("TSMC") – the world's leading manufacturer of sub-10 nm semiconductor chips – noted:

> *"I like the High-NA EUV's capability, but I don't like the sticker price."*

The comment did not reflect a deal-breaking concern, however. One month later, on June 5, it was announced that TSMC had placed an order for the $380 million machine.

One industry analyst summed the current reality: "You can't make leading-edge chips without ASML's machines."

Any chipmaker that doesn't want to get left behind in the never-ending race toward greater computing power has no choice: ASML is the only game in town.

ASML is the only company that can supply the critical lithography machines needed to make today's sub-10 nm computing chips. The parallel computing revolution and AI arms race will likely fuel a surge in demand for these machines in the coming years. Analysts currently expect the company will generate $33 billion in revenue in 2025, up nearly 20% from 2023. Meanwhile, ASML's profit margins are expected to increase to 30%, up from 28% in 2023, as it benefits from growing sales of its latest $380 million High-NA-EUV machines. This translates into a 2025 earnings per share estimate of $24.

The World's Most Important Company

Taiwan Semiconductor Manufacturing (NYSE: TSM) is the world's largest semiconductor manufacturer. While TSM isn't a monopoly in its market like ASML, it's a close second – it produces 90% of the world's most advanced sub-10 nm computing chips.

However, TSM is arguably the most fragile link in the global semiconductor supply chain, because China views the self-ruled island nation of Taiwan as a breakaway province. The Communist Chinese government believes it and the 24 million people who live there will eventually be reunified with mainland China, under the threat of force if necessary. Adding to the uneasiness is the fact that the U.S. has vowed to support Taiwan's claim to independence.

This has created a tense war of words between the U.S. and China over Taiwan's fate as a sovereign nation. A geopolitical showdown over the small island country would have devastating consequences for the world's supply of today's most critical resource: advanced semiconductors.

Bloomberg estimates that a war over Taiwan could cause a $10 trillion hit to the global economy – five times the $2 trillion in losses from the Great Financial Crisis.

Fully appreciating the critical role TSM plays in the global economy starts with understanding why companies like Nvidia, **featured in a previous Big Secret issue**, and virtually all other leading semiconductor companies, don't actually manufacture the chips they sell. Instead, they create the blueprints, and submit those designs to contract manufacturers, like Taiwan Semiconductor.

This outsourcing model made good financial sense when it first emerged in the 1980s. Back then, there were dozens of contract manufacturers, and no one company held a dominant role in the industry.

But as we'll explain here, TSM perfected this outsourcing model better than anyone else over the course of nearly four decades. It's risen from a second-rate manufacturer of low-margin, commoditized chips to the most dominant semiconductor maker in the world – by a wide margin.

With control over 90% of the supply of advanced semiconductors, TSM represents a bigger geopolitical choke point – on a product that's arguably more important to the global economy than oil – than all of OPEC, which controls just 40% of the global oil supply.

The story of how the world's chip supply became dominated by one company all begins with a high-ranking executive that was passed over for a promotion in 1985.

Trial and Error and Lower Costs

By all rights, Morris Chang was a lock to become the next CEO of Texas Instruments ("TI"), America's leading chipmaker, in the mid-1980s.

After immigrating to the U.S. from China and earning a degree in mechanical engineering from MIT, Chang joined TI in 1958. At the time, TI was at the cutting edge of developing the new class of silicon-based semiconductor chips.

Chang proved himself as a brilliant engineer and problem solver early on at TI. In one of his first major assignments, he was tasked with solving one of the company's biggest challenges: fulfilling a contract TI had with IBM to help the computer company manufacture its first silicon-based chips for mainframes.

IBM had problems manufacturing the chips in-house, achieving production yields of just 10% (meaning 90% of the chips produced were thrown away). When TI began producing the same chips for IBM, its yields were even worse. Within four months of taking over, Chang achieved a 20% production yield – or twice as good as what IBM had been able to do.

Chang proved himself as a brilliant businessman as well. By 1967, he had become general manager of TI's burgeoning semiconductor business. In this role, Chang made a key breakthrough that would ultimately set the stage for TSM's dominant business model today.

Chang realized that the biggest economic force governing chipmaking was scale. Chipmaking facilities aren't the typical, commoditized factories that anyone can build. They require immense precision, cutting-edge technology, and – most importantly – trial and error to get yields up (as Chang had done with the IBM chip line). None of this comes cheap.

In the 1960s, industry practice was to pass all of these upfront trial-and-error costs onto the customer. This meant charging high initial prices on new chips, with the aim of reducing prices after recouping their initial investment.

But this presented a catch-22. The high initial prices meant the buyers of chips couldn't afford to order in bulk. Thus, chipmakers might get the prices they wanted, but not in the volume they needed. This meant chipmakers had to run their plants below capacity, resulting in bloated overhead costs that limited their overall returns.

Chang hired Boston Consulting Group ("BCG") to conduct a wide-ranging study on the economics of chipmaking. Chang and BCG came up with a concept called "learning-curve pricing." This was a controversial strategy that involved charging low prices from the beginning, and then consistently reducing prices more – even when the market did not demand it.

Many in the industry thought this was a crazy strategy. And in the short run, they were right. It was very expensive, requiring TI to operate at a loss in the early phase of ramping up new production lines. But Chang had a long-term vision. By cutting prices aggressively, TI was able to secure more business, overcoming the volume hurdles that plagued the traditional industry approach.

By aggressively lowering prices, TI successfully pulled business away from its higher-cost competitors. This strategy turned TI into the industry's largest semiconductor manufacturer. And with all of this volume, TI could spread its fixed costs over a larger number of units, also making it the most profitable chipmaker in the world.

By the late 1970s, Chang's learning-curve pricing strategy had helped TI achieve the critical economies of scale required for profitable semiconductor manufacturing. Chang (along with pretty much everyone else at TI) expected he could eventually become CEO.

But in 1985, for reasons still unclear today, Chang was passed over for promotion to the Texas Instruments CEO role. Chang resigned in frustration. This ultimately created the catalyst for him to accept an invitation to take his talents to the small island country of Taiwan, just off the coast of mainland China, where he used his learnings from TI to upend the global order in semiconductors.

A Chip Manufacturing Giant Is Born

In the 1980s, Li Kwoh-ting (today known as the "father of Taiwan's economic miracle") had a grand vision to transform the country into a global tech giant. In 1987, the economist and political leader recruited Morris Chang to become CEO of a new public-private partnership called Taiwan Semiconductor Manufacturing Company. The goal was to make TSM into the preeminent manufacturer of semiconductor chips.

At the time, Taiwan had no expertise in chipmaking, or any other advanced technology for that matter. The country's manufacturing base centered around low-margin, commoditized products.

What Taiwan did have was a vast supply of cheap labor. This was enough to get TSM started by mass producing the low-margin, older-generation chips that other companies didn't want to make. While TSM's early profits were slim, Chang had learned from his TI experience that long-run value came with greater economies of scale. So he aggressively reinvested into R&D, advanced manufacturing equipment, and on-the-job training for TSM's workforce.

By the mid-1990s, TSM's manufacturing capabilities began closing the gap with some of the world's most advanced chipmakers. And the timing couldn't have been better. Under growing pressure from overseas competitors, U.S. chipmakers began outsourcing production of even their most-cutting-edge semiconductors – opening up a massive source of new and increasingly profitable business for TSM.

Japanese Competition Forces U.S. Chipmakers Into Outsourcing

When the semiconductor industry emerged in America in the 1960s, most chipmakers were fully vertically integrated – meaning they controlled everything from design to manufacturing. When they did outsource production, it was limited to their lower-margin chips.

However, fierce competition from Japanese chipmakers changed everything, starting in the 1970s and accelerating throughout the 1980s. Japanese manufacturers gained an edge with the help of the Japanese government, which provided large tax incentives and funding programs to spur R&D. As a result, total semiconductor R&D went from 2% of Japan's industry total in the early 1970s to 26% by 1977.

By the 1980s, Japanese chipmakers began dominating the market for products like VCRs, which at the time contained the most semiconductors of any consumer device. They also introduced new cutting-edge consumer technology, like the Sony Walkman, a portable music player that was the precursor to Apple's iPod.

By 1986, Japan overtook America as the dominant force in semiconductor manufacturing, with over 50% of the global market share. As a result, U.S. chipmakers had less money to invest into R&D, putting them even further behind, and creating a self-reinforcing downward spiral.

One by one, U.S. chipmakers realized they had only one good option to regain their competitive foothold: outsource manufacturing. This concept gave rise to what we now call the "fabless" semiconductor model. In this model, semiconductor companies focus entirely on researching and designing cutting-edge chips and then sending those blueprints to third parties for mass production – either to other vertically-integrated chip makers with excess production capacity, or to "foundries" that focused exclusively on manufacturing for other companies, but not designing their own chips. In this way, chipmakers still retained the high profit margins from selling cutting-edge semiconductors, while eliminating the huge capital requirements of running a foundry.

For the old guard who grew up in the vertically integrated model, outsourcing production was anathema to their DNA. Jerry Sanders, the flamboyant founder and CEO of Advanced Micron Devices ("AMD") during the 1980s, famously declared that "Real men have fabs," in a less-than-subtle jab at his outsourcing competitors.

But the financial force of capital efficiency couldn't be ignored. By the mid-1990s, a growing number of U.S. chipmakers began adopting the fabless model. Over time, this freed up their capital to invest more into R&D, enabling America's top chipmakers to regain their competitive lead designing the world's most innovative chips – without manufacturing anything themselves.

So even as America's share of chip manufacturing dropped to an all-time low by the end of the 1990s, its share of total chip sales rose above 50% – recovering from a drop below 40% in the 1980s. Meanwhile, Japanese chipmakers fell by the wayside, entering the new millennium with just 29% global market share:

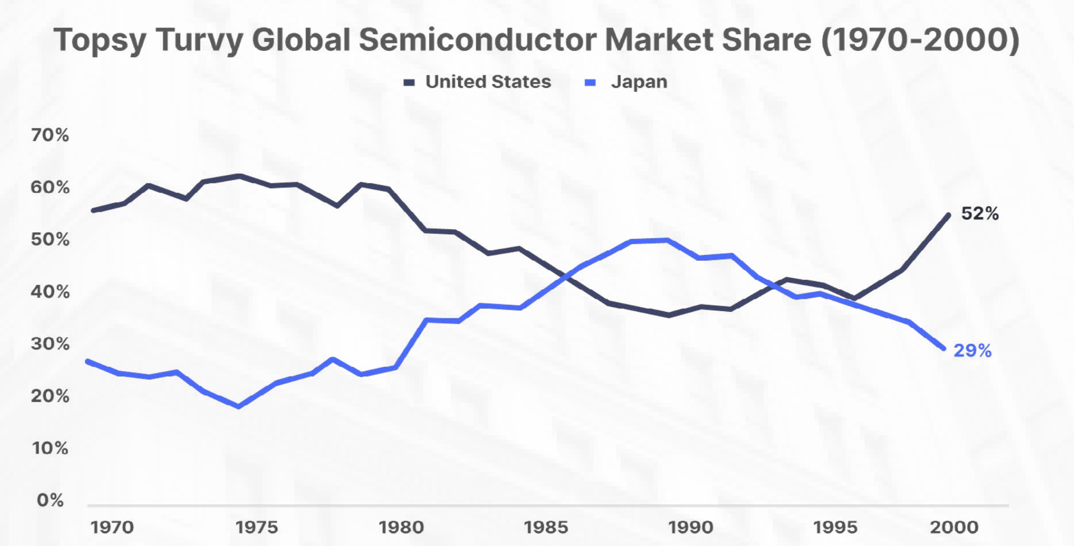

Source: Handotai Sangyo Kenkyusho, 2000

TSM found itself perfectly positioned during this 1990s outsourcing boom.

The company enjoyed a couple of key advantages as the world's first – and for a long time, the only – pure-play semiconductor manufacturer. Before TSM, the only outsourcing options involved partnering with a fully-integrated chipmaker that would effectively rent out their excess production capacity.

This model created inherent conflicts of interest, given that the manufacturer was often in competition with its customers. And once a company learns how to produce a competitor's chips, it's not a giant leap for that company to find ways of mimicking that design.

The brilliance of Chang's approach at TSM was to build a company that purposefully avoided competing with its customers.

TSM's business model was particularly appealing to new startups, which often had valuable IP, but few resources to shield that IP from larger competitors. Thus, TSM became a go-to choice for many up-and-coming fabless chipmakers in the 1990s. One of these startups was a then-little company called Nvidia. Founded in 1993, Nvidia got off the ground with just $40,000 in the bank and $20 million in venture-capital financing. TSM was the key that enabled new entrants like this small California chip designer to enter the semiconductor market, without facing the daunting capital requirements of building its own production facilities.

TSM became Nvidia's primary chip manufacturer in the mid-1990s, and the partnership has flourished ever since. As Nvidia rose to become one of the world's most valuable companies, eaching a market capitalization of over $3 trillion earlier in 2024, it brought TSM along for the ride.

And this long-running partnership underscores another major competitive advantage that sets TSM apart from the competition. The brilliance of TSM's pure-play foundry model meant that the company's fortunes were never tied to any single chip architecture. Instead, TSM's manufacturing capacity followed the demand trends from its most successful customers, like Nvidia, Apple, and Qualcomm.

And it's paid off handsomely. Since going public in 1994, the company has grown its revenue at a CAGR of 18%. Over this 30-year stretch, TSM has only posted three years of revenue declines – each caused by broader macroeconomic slowdowns. These include the 2001 dot-com bust, the 2008 Great Financial Crisis, and the 2023 post-pandemic slump in smartphone and PC demand. Revenue is on track to rebound to new highs this year from the booming demand for AI chips:

Source: Bloomberg

Along the way, TSM aggressively reinvested its growing revenue stream into advancing its production capabilities.

So while TSM started off as a low-margin producer of older-generation chips in the 1980s, everything changed when the outsourcing boom took hold in the 1990s. By the early 2000s, TSM had advanced its capabilities to become one of just 31 companies in the world capable of producing the most advanced 150-nm chips of the day.

Over time, the number of semiconductor manufacturers capable of producing the most advanced chips began dropping off. Part of this was a result of the ongoing trend of outsourcing. Even longtime holdout AMD relented and spun off its manufacturing operation into its own business, GlobalFoundries, in 2009.

And for the companies that continued to both design and make their chips, their foundries were tied in with whatever chip architecture they specialized in. Good or bad, they had to continue producing the chips they designed. As technological trends shifted, many of these companies became cut out from the leading edge of semiconductor manufacturing.

Since TSM had no allegiance to any single chip architecture, it continued grabbing more market share as other foundries fell by the wayside. The never-ending disruptive force of technological evolution became a perpetual boon to TSM, while gradually eliminating other top foundries. As a result, the market for cutting-edge chip manufacturing naturally accrued to TSM.

From 2000 to 2020, the number of companies capable of producing the world's most advanced semiconductor chips fell by 90% from 31 down to just three:

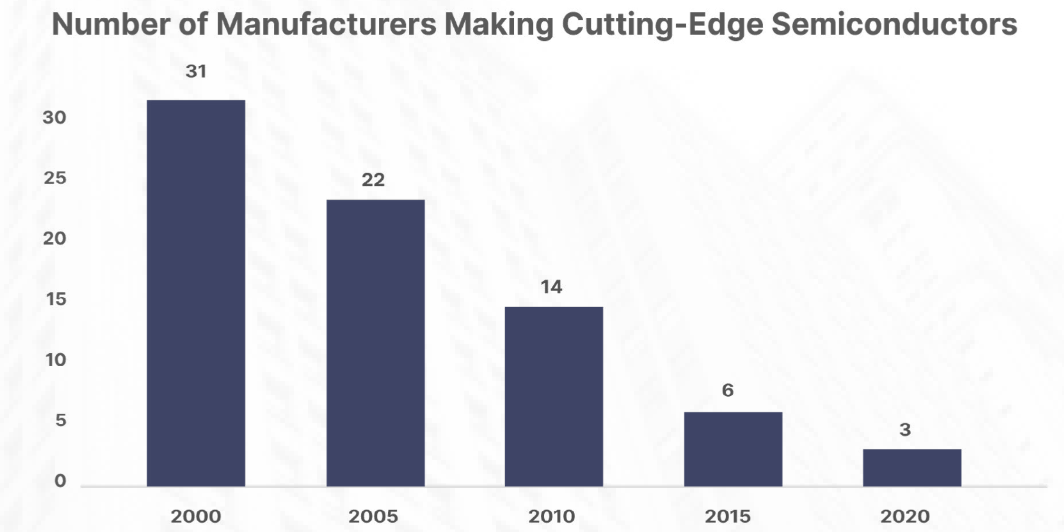

Of those three, TSM is number one. Today, it is the dominant global manufacturer of advanced sub-10 nm chips, with over 90% market share. The other two companies vying for second best include the U.S.-based Intel and Korea's Samsung.

Samsung and Intel are both vertically-integrated manufacturers that primarily make chips for their own product lines. They both offer foundry services for other chipmakers, but this is a minor subsegment of their overall businesses. Moreover, Intel and Samsung both face the same constraint discussed earlier: they often compete head-to-head with their foundry customers.

TSM, on the other hand, is in a league of its own. It's the world's only pure-play foundry with sub-10 nm production capacity that dedicates 100% of its focus to supplying customers, rather than competing against them with its own designs.

Thus, for many of the world's largest fabless semiconductor companies, TSM is the only game in town. It has secured long-standing partnerships to supply advanced chips for virtually all of the world's largest semiconductor companies. Its key customers include Apple, Nvidia, Qualcomm, Broadcom, and AMD.

TSM enables these companies to generate enviable returns on capital, by handling the capital-intensive business of manufacturing the chips these tech giants design. Every year, TSM spends $30 billion on capital expenditures – or roughly double that of the five largest fabless chip makers combined:

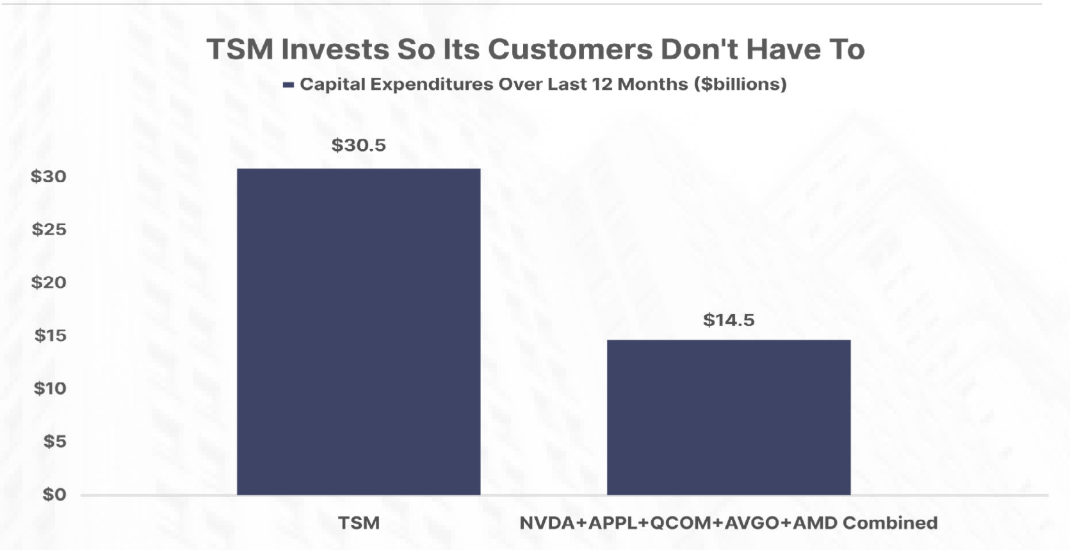

Aside from Intel and Samsung, no other company in the world invests in semiconductor manufacturing capacity at the scale of TSM. And that's one of its most powerful competitive advantages today.

In order for any company to displace TSM's dominant market position, it would first need to invest tens of billions of dollars into state-of-the-art production facilities. But that's only the start. They would also have to spend many years mastering the art and science of optimizing their production yields.

TSM's semiconductor plants are some of the most precise manufacturing operations on the planet. The sub-10-nm patterns TSM etches into silicon wafers are so intricate that a single speck of dust can ruin an entire batch of silicon wafers. This requires elaborate "cleanrooms" that are 10,000 times more sterile than the operating rooms in a modern-day hospital. Engineers must at all times work in protective suits, masks, and gloves. The air is so sterile and free of moisture that engineers must consume gallons of water each day to stay hydrated.

Ticker	Description	Latest Close	Initial Analysis	Yield	Status
Default Group					
ABNB AIRBNB-A	Home Sharing Marketplace	$146.88	05/31/2024	N/A	Buy Under $100
SHW Sherwin-Williams	Specialty Chemicals	$325.46	04/14/2023	0.88%	Buy Under $150
HD Home Depot	Home Products Stores	$366.08	11/10/2023	2.46%	Buy Under $240
PAYC Paycom Software	Application Software	$154.79	01/26/2024	0.97%	Watchlist
Parallel-Processing Revolution					
ASML ASML Hold NY Sp ADR	Semiconductor Supplier	$924.15	06/14/2024	N/A	Watchlist
TSM Taiwan Semi Sp ADR	Semiconductor Manufacturer	$171.87	06/14/2024	1.26%	Watchlist
NVDA NVIDIA	Semiconductor Designs	$121.09	06/07/2024	0.03%	Watchlist
ARM	Semiconductor Designs	$158.33	06/28/2024	N/A	Watchlist

Despite the elaborate protections, speed is of utmost necessity. For that, TSM has installed some of the most sophisticated, high-speed robotics equipment in the world. From a mounted track in the ceiling of their plants, robotic shuttles transport chips through various stages of the assembly process, some of which are worth their weight in diamonds. The robotic shuttles in TSM's plant travel a quarter-million miles each day, equal to 10 trips around the planet.

Recall from earlier, Morris Chang first rose to prominence at Texas Instruments by figuring out how to boost production yields up to 20% in the early 1960s, which was considered a good yield at the time. During the era of U.S. dominance in semiconductor manufacturing, the most efficient chipmakers achieved yields of 50%. When Japanese companies began displacing their American counterparts in the 1980s, they reached an average production yield of 60%. Today, Taiwan Semiconductor has reached 80% production yields – the best in the world. But it took many years of painstaking trial and error to get there.

As a Taiwan-based reporter for BBC recently explained, the "secret sauce" behind TSM's manufacturing prowess is equal parts art and science:

> *"Making microchips is engineering. But it's also more than that. Some have likened it to cooking – like a gourmet feast. Give two chefs the same recipe and ingredients – the better cook will make the better dish."*

Over its more than four-decade history, TSM has developed an irreplaceable body of institutional knowledge around chip making. By continuously tweaking its operating procedures through a never-ending series of trial and error, the company has perfected both the art and the science of chip making. While a competitor could try to replicate this process, it would take a herculean effort of time and human resources, and mountains of capital.

This presents a similar catch-22 situation that protects ASML's secure position. In order to pull business away from TSM, a competitor would first need to replicate TSM's incomparable scale and efficiency. Otherwise, a competitor could never compete effectively on price. But in order to match TSM's scale and efficiency, a competitor would first need to attract enough business to justify the gargantuan investment in time and money.

The only viable solution would involve substantial external funding from a government. Even then, success is far from guaranteed. We're seeing a real-world case study of this playing out in the U.S. today.

As we mentioned earlier, in 2022, the U.S. government passed the CHIPS Act, which authorized $52 billion in funding and tax incentives for the construction of new U.S. semiconductor manufacturing plants. This is part of the Biden administration's efforts to reduce America's reliance on Taiwan for today's most advanced chips. Along those lines, the administration persuaded TSM to commit to building a $40 billion state-of-the-art plant in Arizona.

That $40 billion price tag would make it among the most costly manufacturing plants in the world, explaining why most chips are made overseas. America has become an expensive place to build large-scale manufacturing facilities. The price of land, materials, and (most important) labor is much higher compared with other key chip-making regions, like Korea and Taiwan. Consider that the $75,000 average household income in the U.S. is roughly five times that of the $16,000 in Taiwan.

TSM broke ground on the facility in December 2022, generating much fanfare about the coming resurgence of U.S. domestic semiconductor manufacturing. But the initial enthusiasm has now given way to a bleak reality.

TSM has struggled to recruit enough skilled workers even to complete the initial construction of the factory – getting significant pushback from labor unions. Meanwhile, TSM has also battled the endless red tape and bloated costs associated with the American bureaucratic machine. As Elon Musk recently explained, large-scale construction projects have effectively been made illegal in Western economies:

> *"In the West, I think we have created regulatory gridlock where almost everything is illegal. This is why they can't build a high-speed rail in California. They spent $7 billion and there's a 1600-foot section, it's all they have to show for it and it doesn't even have rails on it."*

Production at TSM's Arizona plant was originally slated to begin this year, but it's since been delayed to 2025 because construction is not yet complete. Even when the factory comes online, many are skeptical that the plant will be able to compete on the global scale given the high costs of manufacturing in the U.S. One skeptic is a former TSM chairman, who described the Arizona chip plant as an "expensive, wasteful exercise in futility." He also estimated that the U.S.-made chips will cost 50% more than those produced in Taiwan.

While the idea of a vibrant U.S. semiconductor manufacturing industry is appealing, we remain doubtful that the math will work – absent a radical overhaul in U.S. industrial policy. As California's high-speed rail system shows, simply throwing money at the problem is not a solution. U.S. policymakers must dramatically reduce the bureaucratic burden on manufacturers to give them a chance on the global stage. But Washington isn't even trying.

The bottom line: despite more than $50 billion in direct support from the U.S. government, plus tens of billions more from the private sector, all signs indicate TSM will retain its dominance over the global chip supply for years to come. The situation reminds us of Warren Buffett's famous gut check for an enduring competitive moat, which he once described in reference to his favorite soft-drink maker, Coca-Cola (KO):

> *"If you gave me $100 billion and said take away the soft drink leadership of Coca-Cola in the world, I'd give it back to you and say it can't be done."*

Like Coca-Cola, TSM's wide competitive moat means even the most well-funded competitors can't displace its dominant market position. As the only game in town for advanced semiconductor manufacturing, TSM commands premium pricing power for its chips. This is evident in the fact that the company makes more money on its chips than do its largest customers, most of which earn double-digit profit margins themselves. Generating an average of nearly 40% net income margins over the last three years, TSM is one of the most profitable companies in the world:

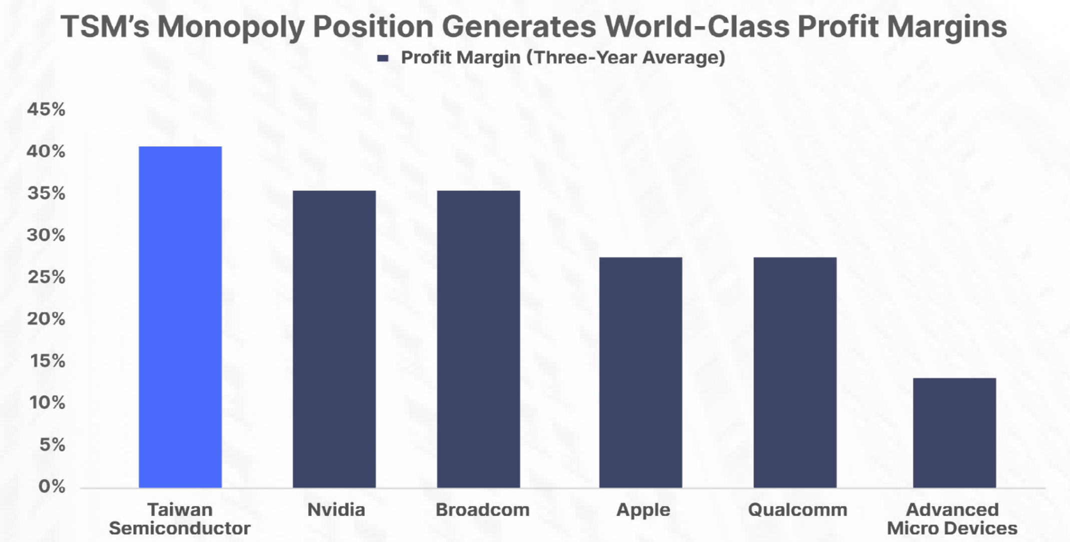

Porter & Co.

THE BIG SECRET ON WALL STREET

The Missing Link

- The CPUs Powering Today's Parallel-Processing Boom
- Transforming Data Centers Into Supercomputers

FROM THE DESK OF PORTER STANSBERRY

The Missing Link

The CPUs Powering Today's Parallel-Processing Boom

Transforming Data Centers Into Supercomputers

No one insulted Sir Alan Sugar's pet robot and got away with it.

Especially not cheeky technology journalist Charles Arthur… who'd dared to suggest in a 2001 article in British online newspaper The Independent that Lord Sugar's latest invention was a "techno-flop"… difficult to operate and even harder to sell.

(Lord Sugar with his pet robot)

Sugar's brainchild, the E-M@iler – a clunky landline phone with attached computer screen that allowed you to check your email – wasn't flying off the shelves of British electronics retailers. With smartphones still a decade or so away, the idea of browsing emails on your telephone just seemed bizarre.

And the E-M@iler – despite its appealing, friendly-robot presence – had plenty of drawbacks: users were charged per minute… per email… and were bombarded by non-stop ads on the tiny screen. Customers' phones were auto-billed each night when the E-M@iler downloaded the day's mail from the server. (By the time you'd owned the phone for a year, you'd paid for it a second time in hidden fees and charges.)

But scrappy Sir Alan, who'd grown up in London's gritty East End and now owned England's top computer manufacturer, Amstrad, wasn't about to admit that he'd pumped 6 million pounds, and counting, into a misfire.

Sugar wasn't afraid to fight dirty. And he had a weapon that Charles Arthur, the technology critic, didn't: 95,000 email addresses, each one attached to an active E-M@iler.

On an April morning in 2001, Lord Sugar pinged his 95,000 robots with a mass email. It was a call to arms.

> *I'm sure you are all as happy with your e-mailer as I am," he wrote. "The other day… the technology editor of the Independent said that our e-mailer was a techno-flop… It occurred to me that I should send an email to Mr Charles Arthur telling him what a load of twaddle he is talking. If you feel the same as me and really love your e-mailer, why don't you let him know your feelings by sending him an email.*

In a move that would likely get him in legal trouble today, Sugar included the offending critic's email address… then sat back and waited.

A week or so later, the unsuspecting Charles Arthur returned from vacation to find his inbox bursting with 1,390 messages.

But – in a surprise turn of events for Lord Sugar – most of the missives were far from glowing endorsements of the E-M@iler…

Instead, scores of ticked-off robot owners took the opportunity to complain to the press about their dissatisfaction with their purchase. *"This is by far the worst product I have bought ever." "I won't be getting another one, I think they're crap." "This emailer I am testing is the second one and I now know that it performs no better than the first one." "I WANTED TO SAY THAT I'M NOT HAPPY WITH MINE!!!"* And so on.

The whole story – including the flood of bad reviews – made the news and made Lord Sugar, probably rightly, out to be an ass.

And the bad news kept piling up for him. Over the next five years, the E-M@iler continued to disappoint. In 2006, the price was cut from £80 to £19, with Amstrad making a loss on every unit. After a string of additional financial and operational business misfires, Lord Sugar sold Amstrad to British Sky Broadcasting for a fire-sale price of £125 million… a far cry from its peak £1.3 billion valuation in the '80s.

Fortunately for Sir Alan Sugar, he found a fulfilling second career post-Amstrad. Today, he serves as Britain's answer to Donald Trump, yelling "You're fired!" at contestants on the long-running English version of the TV show The Apprentice.

In a way, though, the E-M@iler fired Lord Sugar first.

It was an ignominious finish for a company that, in the 1980s, had dominated 60% of the market share for home computers in England.

And in the end, Sugar's most lasting contribution to computer technology wasn't even something he created directly. Today, we have the super-computer chip that's the "brains" of the Parallel Processing Revolution… largely because Lord Sugar got in a fight.

A Call to ARM1

From his childhood days hawking soda bottles on the streets of London, Alan Sugar was a hustler. He launched his technology company, Amstrad (short for Alan Michael Sugar Trading) selling TV antennas out of the back of a van in 1968, then graduated to car stereos in the '70s and personal computers in the '80s.

With the 1980s personal computer boom – as technology advanced, and formerly massive mainframes shrunk down to desktop size – came fierce competition....

On the American side of the pond, Apple and Microsoft jockeyed for the pole position (a dance that still continues today). The English home-computer arms race wasn't as well known outside of trade publications… but in the early '80s, British nerds would have drawn swords over the burning question of "Acorn or Sinclair."

Acorn (known as the "Apple of Britain") was the more sophisticated of the two, with a long-standing contract to make educational computers for the BBC (British Broadcasting Company), while Sinclair traded on mass appeal, producing Britain's best-selling personal computer in 1982. (Their feud was dramatized in a 2009 BBC documentary called The Micro Men.)

Into this fray waded – you guessed it — Sir Alan Sugar, who knew nothing about programming or software but was determined to propel his protean tech company, Amstrad, to the top of the industry. He did this in the simplest way possible: He bought the more popular combatant, Sinclair, wholesale in 1986.

Just as he'd hoped, the existing Sinclair product line — plus a successful monitor-keyboard-printer combo designed by Sugar himself – propelled Sugar's company to a £1.3 billion valuation by the mid-'80s.

Tiny £135 million Acorn, left out in the cold, couldn't compete. Or could it?

Acorn fought valiantly before selling out to an Italian firm and ultimately bowing out of the computer biz in the early '90s. And out of that effort grew the super-speedy chip that, today, powers the Parallel Processing Revolution.

Acorn's proprietary chip, called ARM1, was based on a "reduced instruction set computer" architecture, or RISC. It was 10 times faster than the "complex instruction set computer" (CISC) CPU chips in Lord Sugar's computers – and it could be manufactured more cheaply, too. The ARM1-powered computer retailed at one-third the price point of a mainstream PC.

But speedy or not, Acorn's specialized computers couldn't best Lord Sugar's mass-produced PCs. Unlike Sugar's IBM-compatible computers, Acorn's machines couldn't run the popular Microsoft Windows operating system.

Without Windows – and without Lord Sugar's flair for headline-grabbing chaos — the ARM1-powered computer launched with little fanfare, selling just a few hundred thousand units over the next several years.

But long after Lord Sugar boxed up his E-M@ilers and joined the set of The Apprentice… and long after Acorn cashed in its chips and closed its doors… the ARM1 chip has survived and thrived, thanks to a spinoff partnership with Apple.

And ultimately, that tiny piece of silicon developed into the "brain" that drives Nvidia's supercomputing architecture… and powers the Parallel Processing Revolution today.

Remarkably – like the GPU (graphic processing unit) monopoly we explored in "The Big Bang That No One Noticed" – the intellectual property to this chip technology is controlled today by just one company: **ARM Holdings (Nasdaq: ARM)**. While it doesn't get much attention in the media, ARM is poised to become one of the biggest winners from today's parallel-computing revolution.

The Company That Ties It All Together

In our The Big Bang That No One Noticed", which kicked off our Parallel Processing Revolution series, we explained how Nvidia is far more than just an artificial intelligence ("AI") chipmaker. Over the last two decades, the company has laid the foundation for a technological revolution – one that's now changing the very concept of what a computer is.

The combination of Nvidia's super-powered GPUs that greatly increase the processing speed of computing, and its CUDA (Compute Unified Device Architecture) software network, unlocked the parallel-processing capacity required for training the large language models ("LLM") powering today's AI revolution. But that was just the start. The massive computational workloads of training LLMs – which scan large amounts of data to generate simplified, human-like text – presented a new challenge beyond computing speeds. The process of training LLMs to recognize patterns across huge swaths of data created an explosion in memory demand.

Consider the number of data points, known as parameters, used to train today's cutting-edge LLMs like ChatGPT. Training the first GPT-1 model in 2018 required 120 million parameters. Each new iteration required an exponential increase in training data. ChatGPT-2 was trained on 1.5 billion parameters in 2019, followed by 175 billion for ChatGPT-3 in 2020. The number of parameters in the latest iteration, GTP-4, hasn't been disclosed, but experts estimate it was trained on approximately 1.7 trillion parameters – a 100x increase in a matter of a few years.

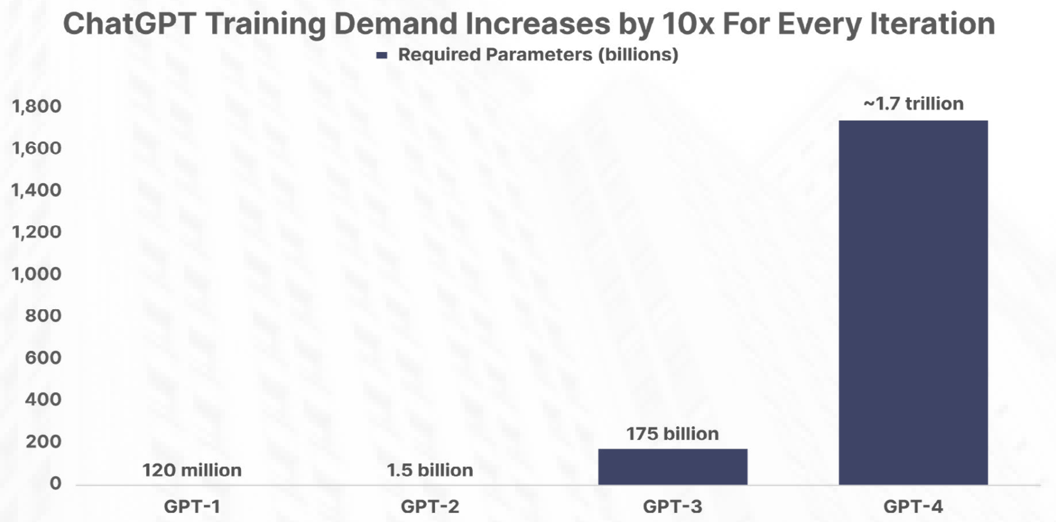

LLMs must hold mountains of data in memory while training on these parameters. In computer science, memory capacity is measured in terms of bits and bytes. A bit refers to the smallest data unit that references a single 1 or 0 in binary code. A byte refers to an 8-bit data structure used to reference a single character in memory, such as a letter or number. One trillion bytes make up 1 terabyte – and hundreds of terabytes are required for training today's most advanced LLMs.

The challenge: even cutting-edge GPUs, like Nvidia's H100, hold less than 10% of the amount of memory needed for training today's LLMs – only about 80 gigabytes (one gigabyte equals roughly 1.1 million bytes) of memory are inside each individual H100 chip.

In the data-center architectures from just 2019, to get more memory, systems engineers would link multiple chips together via ethernet cables. This hack was sufficient to handle most data-center workloads before the age of AI.

But things have changed dramatically since then. Now, hundreds of terabytes of memory storage and transmission capacity are needed to power data centers. Almost no one in the industry anticipated that level of change would happen so quickly – but Jensen Huang did. As far back as 2019, the Nvidia CEO foresaw the future need to connect not just a few chips, but hundreds of chips in a data center. Huang reimagined the data center from a series of compartmentalized chips working independently on different tasks, to a fully-integrated supercomputer, where each chip could contribute its memory and processing power toward a single goal – training and running LLMs.

In 2019, only one company in the world had the high-performance cables capable of connecting hundreds of high-powered data-center chips together: Mellanox, the sole producer of the Infiniband cables (mentioned in The Big Bang That No One Noticed"). That year, Nvidia announced its $6.9 billion acquisition of the networking products company. During a conference call with journalists discussing the acquisition, CEO Huang laid out his vision for this new computing architecture:

> *Hyperscale data centers were really created to provision services and lightweight computing to billions of people. But over the past several years, the emergence of artificial intelligence and machine learning and data analytics has put so much load on the data centers, and the reason is that the data size and the compute size is so great that it doesn't fit on one computer... All of those conversations lead to the same place, and that is a future where the datacenter is a giant compute engine... In the long term, I think we have the ability to create data-center-scale computing architectures.*

David Rosenthal, host of the tech podcast *Acquired*, called Nvidia's purchase of Mellanox "one of the best acquisitions of all time." It provided the missing link the chipmaker needed to harness the power of hundreds of data-center chips together into a massive and explosively fast architecture. Less than four years after the acquisition, Huang's supercomputing vision became a reality, in the form of the Grace Hopper Superchip architecture, which combines the Hopper GPU (based on Nvidia's H100 chip) and the Grace CPU (more on this below).

The big breakthrough in combining these was Nvidia's ability to package Mellanox Infiniband technology into its proprietary NVLink data-transmission cables. The NVLink system in the Grace Hopper architecture transmits memory data nine times faster than traditional ethernet cables. This enabled Nvidia to connect up to 256 individual Grace Hopper chips together, and tap into the full memory bank of both the Hopper GPU and the Grace CPU.

It was a major step in the parallel processing revolution.

On to the Next Step in the Revolution

The end result: Nvidia transformed its already-powerful data-center GPUs into a supercomputer, with a total addressable memory bank of 150 terabytes. A number that is frankly too large to comprehend, 150 terabytes is 150 trillion bytes of data storage capacity, or nearly 2,000 times the memory capacity of a single H100 GPU, and about 10,000 times that of a typical laptop hard drive.

Therein lies the big breakthrough that has allowed Nvidia to upend the concept of modern-day computing. And no other company has come close to putting together this full ecosystem for transforming the data center into a supercomputer.

Today, Nvidia owns virtually all of the hardware, software, and networking technology necessary to transmit terabytes of data at lightning quick speeds throughout the data center. Chamath Palihapitiya, one of Silicon Valley's most high-profile venture capitalists, **described** the significance of the Grace Hopper architecture on the All-In podcast following Nvidia's release of it:

> *This is a really important moment where that Grace Hopper design, which is basically a massive system on chip, this is like them [Nvidia] going for the jugular... that is basically them trying to create an absolute monopoly. If these guys continue to innovate at this scale, you're not going to have any alternatives. And it goes back to what Intel looked like back in the day, which was an absolutely straight-up monopoly. So if Nvidia continues to drive this quickly and continues to execute like this, it's a one and done one-company monopoly in AI.*

However, there's one piece of the puzzle that Nvidia tried, and failed, to corner the market in. That piece is boxed in red in the diagram below – the Grace CPU – the only ingredient of technology in this system that Nvidia relies on another company for:

The Missing Piece of the Supercomputer Puzzle

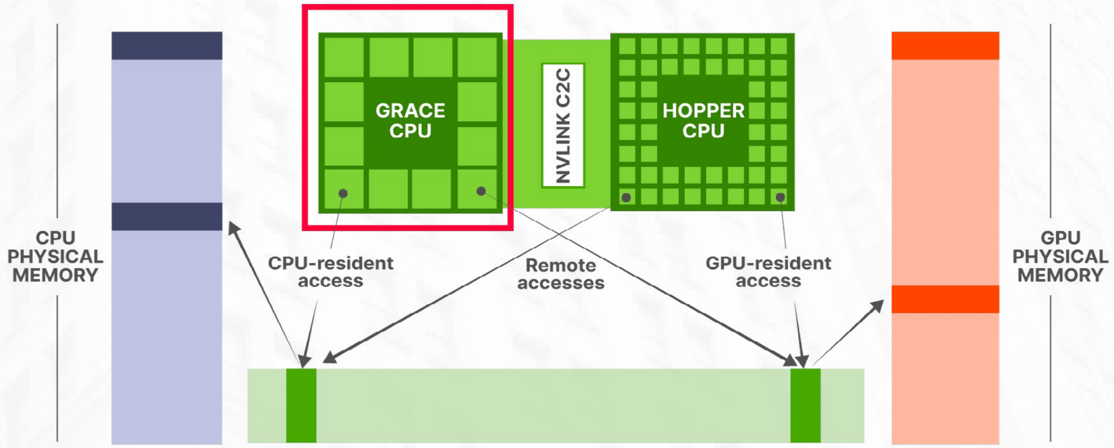

Even in today's new era of parallel computing, there are still some tasks that require running computations in a serial sequence. As a highly simplified analogy, the serial-processing CPU acts like the brains that guide the brawn of the parallel-processing GPU. And just like the human brain, the CPUs in modern-day data centers require huge amounts of energy.

ARM Holdings is the only company in the world that owns the intellectual property underlying the most energy-efficient CPU designs in the market, including the Grace CPU. And while it gets much less attention in the financial press relative to Nvidia's GPU dominance, ARM holds the keys to the critical CPU designs fueling the current AI arms race. In this issue, we'll show why ARM is poised to become one of the biggest winners from today's parallel computing revolution.

A Global Juggernaut Born in a Barn

ARM's CPU technology and business model has come a long way since its formation as a spinoff from Acorn Computers in 1990. Today, the company no longer produces chips. Instead, it licenses its intellectual property to other semiconductor companies that design chips based on ARM's technology. In the last 25 years, over 285 billion semiconductors have been manufactured using ARM-based designs. This includes 28.6 billion chips in 2023 alone.

The technology that ARM licenses is known as the "instruction set architecture" (ISA) for semiconductors. This refers to the set of instructions working behind the scenes of every programmable electronic device that determines how the chip hardware interfaces with the device's software. Think of the ISA as the programming language of the computer chip, which contains a set of commands that a microprocessor will recognize and execute. This is how electronic devices transform the binary language of 1s and 0s into functional outcomes on a given device.

ARM's exclusive focus on designing chip architectures gives it a discrete place in the industry, as one industry analyst describes:

> *Most people think about a device. Then maybe if they're really sophisticated, they think about the chip, but they don't think about the company that came up with the original ideas behind how that chip operates. But once you understand what they [ARM] do, it's absolutely amazing the influence they have.*

That influence includes a 99% market share in global smartphone CPU architectures, along with a major presence in tablets, wearable devices, smart appliances, and automobiles. ARM's next big opportunity will come from its aggressive expansion beyond smartphones into the new data-center CPU architectures powering today's parallel computing revolution.

A key advantage for ARM in today's data-center CPUs is driven by the same feature that allowed its chip architecture to take over the global smartphone market. To understand the source of this advantage, let's go back to ARM's founding when it first broke into the emerging market of handheld devices in the 1990s.

ARM Finds Its Killer App in the Mobile Revolution

In 1990, Apple Computer was three years into developing one of the world's first personal digital assistants (PDA), the Newton. The company originally hired AT&T to develop the Newton's CPU, but after a series of setbacks and cost overruns, it ditched AT&T and began searching for a new partner.

Unlike PCs, which plugged into wall outlets, the main constraint of mobile devices was battery life. One shortcoming of AT&T's CPU was its high energy consumption, which limited the battery life. The chip was based on the industry-standard architecture at the time, known as a "complex instruction set computer" (CISC).

Apple homed in on Acorn's energy-efficient RISC architecture emerging at the time as the perfect fit for the CPU design in its new mobile device.

The key difference between Acorn's RISC architecture and the CISC chip designs boils down to how these two approaches make use of something called the clock cycle in a CPU. The clock cycle is the timing mechanism the CPU clock uses to synchronize operations across various chip components.

In a CISC architecture, the CPU performs many different computing tasks (i.e., transferring data, arithmetic operations, or accessing memory) all from a single instruction that can vary in length and complexity, requiring multiple clock cycles to execute.

Conversely, the RISC architecture limits each instruction set to a simpler, fixed-length format that only takes up a single clock cycle. These simpler instructions performed on a single clock cycle enabled RISC-based chips to run more computations, faster, even with less powerful semiconductor chips versus those based on the CISC architecture.

Leaving the PC market behind, Acorn spun off its division dedicated to ARM CPU technology into a joint venture (JV) with Apple to develop the Newton CPU. Apple contributed $2.7 million to acquire 46% of the JV, and Acorn contributed its technology and 12 employees for another 46%. A third company called VLSI Technology – which previously made Acorn's ARM-based CPU chips – became the manufacturing partner in exchange for the remaining 8% equity stake.

The new entity, formed in 1990, was called Advanced RISC Machines, which later became ARM Holdings.

The ARM design team spent the next three years working closely with Apple engineers to develop a custom CPU chip for the Newton PDA, launched in 1993. Despite achieving all of Apple's CPU requirements for processing speed and efficiency, the product itself was a flop. The critics said that Apple over-engineered the Newton's features, building in unnecessary functionality that consumers didn't want, resulting in a hefty price tag of $700 (about $1,500 in today's dollars).

ARM's CEO at the time, Robin Saxby, learned a lesson from the endeavor and realized that the company couldn't stake its future on single product lines. So he devised what would become an incredibly successful, and highly capital efficient business model that remains in place today: the licensing and royalty model.

Instead of developing chip architectures around individual product lines, ARM would develop platforms of chip architectures, which could be fine tuned based on a specific end market. In this way, ARM could distribute its chip architectures to many different companies, diversifying its risk and tapping into a wider revenue base. And it would charge customers an upfront license fee for using its chip architectures, plus a small royalty (typically 1% to 2% of the total CPU cost) on each chip sold.

Taking ARM to the Next Level

After pivoting to this new model, ARM struck a landmark deal in 1993 with mobile-phone maker Nokia and its chipmaker Texas Instruments. At the time, Nokia was developing the precursors to today's smartphones, implementing quasi-intelligent applications like text, email, calculators, and games. These features required CPUs, and just as with the Newton, energy efficiency.

As part of the Nokia deal, ARM developed what became the ARM7 family of microprocessor architectures. This architecture went into the CPU of the Nokia 6110 GSM mobile phone, first launched in 1998. The highly energy-efficient ARM7 RISC-based CPU gave Nokia the best battery life on the market, with a total talk time of 3.3 hours and a standby time of 160 hours. This compared with competing phones that provided around two hours of talk time and 60 hours of standby time.

The high-memory capacity and processing speeds of the ARM7 CPU also enabled Nokia to hold more data, and include extra features like games that came pre-downloaded with the phone. One of them, Snake, became a viral sensation with over 350 million copies produced, becoming one of the world's first major mobile games.

Becoming a status symbol to own, the Nokia 6110 was one of the most popular phones of the mid-1990s. Nokia and ARM flourished together in the years that followed, with new ARM7 architectures developed alongside new Nokia models. This included the 3210, released in 1999, which became the best-selling mobile phone of the era with an incredible 160 million unit sales.

The successful partnership with Nokia not only provided a massive source of licensing and royalty income for ARM, but it led to widespread adoption of the ARM7 platform among other mobile-phone developers. The architecture was ultimately licensed by 165 different companies, and it's been used in over 10 billion chips.

ARM's RISC-based architecture became the leading global standard in mobile-phone CPUs by the late 1990s. The company went public in 1998, and rising revenue and profits sent its stock price soaring. That same year, Apple shut down the Newton and began selling its shares in ARM to shore up its financials. Apple sold its original $2.7 million stake for nearly $800 million, saving the future smartphone leader from bankruptcy and providing the breathing room to turn around its business.

Fast forward a decade to 2007, and ARM once again partnered with Apple on a much more promising project: the CPU architecture for the world's first iPhone. Apple originally selected Intel to develop an x86-based CPU in the first iPhone. However, Intel rejected the offer over pricing, failing to realize the long-term potential of the smartphone market.

As a result, ARM became the go-to supplier for Apple iPhone chips, as well as virtually every other smartphone on the market. Today, ARM's CPU architectures are found in 99% of global smartphones. Along the way, ARM reinvested its growing stream of profits into other areas suited for its highly energy-efficient designs.

This included the market for digitally-connected internet of things ("IOT") devices ranging from smart home appliances to a variety of blue-tooth enabled devices, to wearable consumer electronics, like the Apple Watch. ARM-based chips make up 65% of the CPUs in the IOT market today. The company also has a growing presence in CPUs for both traditional vehicles and autonomous-enabled vehicles, including the ARM-based chips used in Tesla's Autopilot and Full Self-Driving systems. ARM has a 41% market share in automotive CPUs.

But looking ahead, the most promising market for ARM's RISC-based CPUs lies in the new supercomputing data-center architecture.

The Rise of ARM and Fall of Intel in the Data Center

Starting in the 1990s, Intel began adapting its widely-popular x86 PC architecture into the data center – huge warehouses filled with computers for storing large amounts of data – taking market share from IBM, the industry leader at the time. Intel's key advantage came from its economies of scale. It was already mass-producing x86-based chips for the PC market, and thus it enjoyed a low-cost advantage given the sheer volume of x86 chip manufacturing capacity. Meanwhile, Intel leveraged its dominant research-and-development (R&D) budget from its highly profitable PC chip business into developing newer, more capable x86-based chips for the data-center market.

As a result, Intel's market share in CPUs grew to over 90% by 2017.

ARM's RISC-based CPUs, meanwhile, held a minuscule single-digit market share position in the data center. As recently as 2020, the company wasn't yet focused on this sector because the single biggest advantage of its RISC architecture – energy efficiency – took a backseat to price, favoring Intel's greater economies of scale.

But by 2020, Nvidia CEO Jensen Huang realized that energy demands would explode in the new era of data-center supercomputers. And he also knew that ARM's energy-efficient RISC architecture would become the new gold standard of data-center CPUs in this new paradigm. This was the missing piece of the supercomputing puzzle Nvidia needed in order to control the entire ecosystem of parallel processing.

At the time, Japanese holding company Softbank owned ARM (it bought the company for $32 billion in 2016). In September 2020, Nvidia made a $40 billion offer to acquire ARM from Softbank, which agreed to the sale. However, the U.S. Federal Trade Commission blocked the sale based on antitrust concerns. Softbank maintained full ownership of ARM until September 2023, when it listed 10% of its stake for sale in a public offering.

Short of owning ARM, Nvidia did the next best thing: it worked closely with ARM on developing a CPU optimized for Nvidia's new data-center supercomputer architecture. The effort came to fruition in the form of the Grace CPU, paired with the Hopper GPU in the Grace Hopper architecture mentioned above.

Nvidia and ARM worked together to optimize the RISC-based instruction set in the Grace CPU for both performance and energy efficiency. And with great success.

Compared with competing x86 alternatives, the new design boosted the performance of Nvidia's H100 GPUs by a factor of 30x while reducing energy consumption 25x. The end result: the new design created enough energy savings to run an additional 2 million chatGPT queries versus a comparable x86 CPU chip.

Given that each chatGPT query requires roughly 10x as much energy as a non-AI-powered Google search, this adds up to major cost savings for data-center operators. The immense energy needs of the AI revolution is the key reason why ARM-based CPUs are set to displace Intel's x86 architecture in the data center going forward.

And unlike the first time ARM went head-to-head against Intel during the PC revolution of the 1980s, it's now playing from a position of superior strength. One of the key features of building an enduring competitive moat around computing technologies – whether in software, hardware, or instruction set architectures – is building a robust developer ecosystem around that technology. A developer ecosystem refers to the network of external software and hardware engineers that help advance a technology platform.

When Intel and Microsoft created what was called the Wintel duopoly starting in the 1980s, the dominance of this software-hardware regime led to a robust developer ecosystem that by the late 1990s grew to millions of third-party software and hardware engineers. This ecosystem helped Microsoft develop new software features for advancing its most popular applications, like Excel and Word.

A similar network of developers sprang up to advance the x86 chip architecture. This created a self-reinforcing cycle: the progression of new Windows features and x86 processing capabilities made it the go-to choice for virtually all major PC manufacturers. This, in turn, provided a powerful draw for attracting more developers into the Wintel ecosystem, leading to further technological dominance, and so on.

Now, ARM is taking a page from the same playbook by partnering with the new king of the data center – Nvidia. Over the last several years, Nvidia has begun optimizing its CUDA software platform to mesh with ARM's CPU architecture. This has allowed Nvidia's 4.5-million-strong developer network to begin designing data-center applications around ARM's RISC-based CPU architecture.

But that's just the tip of the iceberg. ARM has also invested heavily into creating its own world-class developer network over the last 35 years. The company now has around 18 million developers who have contributed 1.5 billion hours toward everything from chip designs and operating system integrations to software tools. This puts ARM's developer network at fo ur times the size of Nvidia's, and more than half the size of tech giants Apple and Microsoft.

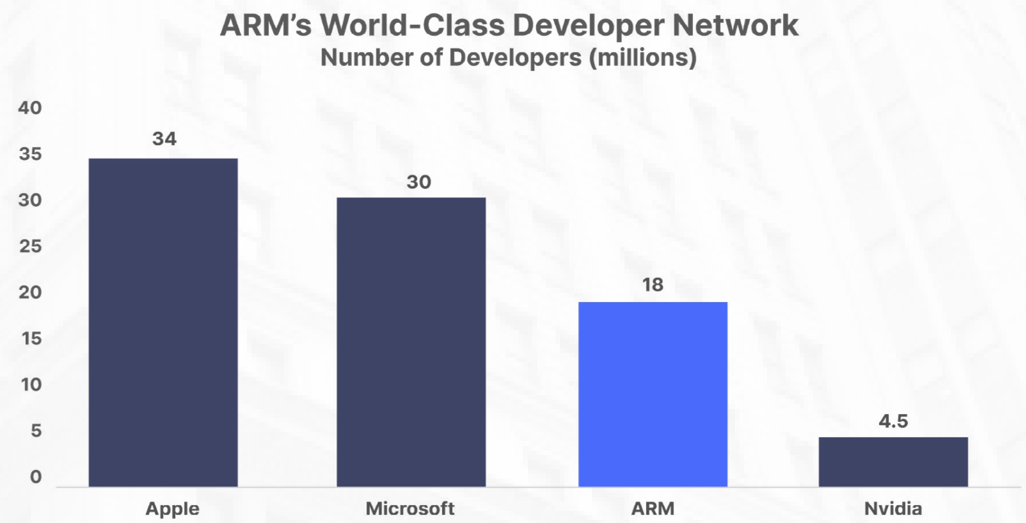

With nearly 30 billion ARM-based chips produced each year, it's the most widely adopted processor architecture in the world. This expanding ecosystem will continue fueling ARM's growing competitive advantage, as the company explains in its financial filings:

> *The breadth of our ecosystem creates a virtuous cycle that benefits our customers and deeply integrates us into the design cycle because it is difficult to create a commercial product or service for a particular end market until all elements of the hardware and supporting software and tools ecosystem are available.*

Beyond this ecosystem, ARM's other competitive advantage lies in its massive intellectual-property portfolio that includes 7,400 issued patents and 2,500 pending patent applications. This is a result of its heavy emphasis on research, with an industry-high 83% of ARM employees working in research, design, and technical innovation.

This has solidified ARM's role as the global leader in RISC-based CPU architectures by a wide margin. And its expertise in this energy-efficient chip design has made it the go-to choice for the world's largest data-center operators. This includes Amazon, Microsoft, and Google – all of whom are racing to create their own data-center chips. Consider a few examples…

In November of 2023, Amazon Web Services ("AWS") announced it would use ARM's Neoverse V2 core inside of its fourth-generation Graviton series of CPUs for its cloud-computing infrastructure. The ARM Neoverse is based on the same ARMV9 architecture used in Nvidia's Grace CPU.

A report from Amazon based on customer testimonies of AWS users shows that the Graviton delivers 20% to 70% cost savings versus comparable x86-based chips from Intel and Advanced Micro Devices ("AMD"). Amazon has become a major partner and customer of ARM, and reportedly makes up 50% of the demand for ARM-based data-center CPUs. Amazon has partnered with ARM since 2018 on the Graviton, and is now working on its fifth-generation ARM-based Graviton chip.

Then in April, Google announced it was developing a custom ARM-based CPU for AI workloads in its data centers. Named Axion, the CPU is already running Google's internal AI workloads for optimizing things like its YouTube ads. It plans to roll out the chips to its business customers later in 2024. Google has said that the Axion chip will deliver 50% better performance than comparable x86 chips produced by Intel.

Microsoft is also using ARM-based designs in its Ampere Altra CPU processors for its Azure cloud-computing engines. The software giant has reported a 50% increase in "price performance" (i.e., performance for comparably priced chips) relative to comparable x86 chips.

The bottom line: ARM's energy-efficient RISC CPU architecture is rapidly gaining market share in today's energy-intensive data centers. By delivering roughly 50% better performance than the traditional x86 CPU architecture, ARM is rapidly disrupting Intel's former dominance over the data center. As a result, ARM's market share in data-center CPUs has more than tripled from just 3% in 2020 to 10% tin 2024, while Intel is losing share:

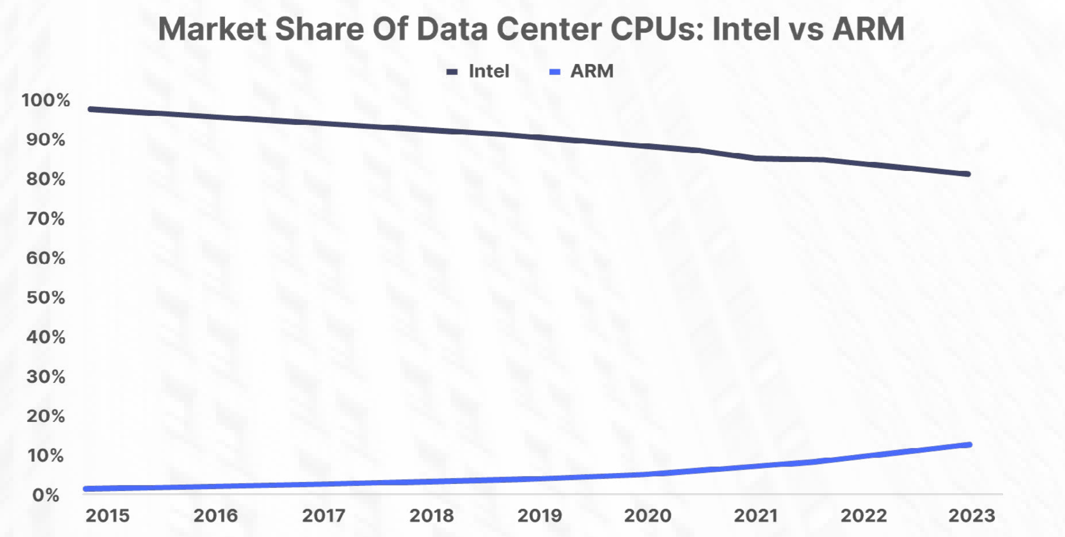

This trend is still in the early innings, and it's expected to accelerate in the years ahead. Industry analysts expect ARM's share of data-center CPUs will double to 22% in 2025, and reach 50% by 2030.

Finally, there's another new source of demand for ARM-based chips in the burgeoning market for AI-enabled personal computers. Unlike previous AI applications, like ChatGPT that run through the cloud, AI computers enable users to run AI applications on their personal computing devices. Analysts estimate that AI-based PCs could make up 40% of global shipments as early as 2025.

ARM is currently working with over half a dozen companies to implement its RISC architecture into chip designs for AI-enabled personal computers. These include Nvidia, Microsoft, Qualcomm, Apple, Lenovo, Samsung, and Huawei.

AI PCs could hold the future for ARM to massively expand its presence into the PC market, given the advantages of its RISC-based architectures over the energy-hungry x86-based chips that currently dominate the PC market.

ARM has already made significant inroads into displacing the x86 architecture in PCs through its long-running partnership with Apple. In 2020, Apple announced it would begin transitioning away from Intel's x86 processors to ARM-based M1

Apple Silicon Chips across its line up of PCs including MacBook Air, MacBook Pro, Mac Mini, and iMac. The M1 chips (and subsequent iterations) have been a massive success for Apple, delivering substantial improvements in performance and battery life.

A Wonderful Business – But at a High Price

ARM generated $3.2 billion in revenue from the 28.6 billion chips produced from ARM-based architectures in 2023 – an increase of 21% year-on-year. On the surface, this revenue figure may seem small relative to the number of units sold – with ARM receiving roughly 12 cents on each chip sold based on its device architecture.

But the key to ARM's business model is that it requires virtually no capital investment, because ARM doesn't manufacture anything. It's purely a technology company that creates the blueprints for chip architectures. It then licenses those blueprints to other companies, like Nvidia, which design their chips based on ARM's architecture. ARM receives an upfront licensing fee, plus a small royalty on every chip sale in perpetuity.

The beauty of this business model is that ARM can create a blueprint once, and sell that same blueprint (or a slightly modified version) to multiple companies. And it collects a perpetual revenue stream for as long as the chip remains in production, which can be for decades. ARM today is earning royalty revenue on chip architectures it designed in the 1990s.

That's the ultimate version of "mailbox money" – consistent incoming revenue that requires no operating expenses or investment, translating into pure free cash flow. This makes ARM one of the most capital efficient businesses in the world, with the company spending less than $100 million in capex each year to produce over $3 billion in revenue. As a result, ARM generates roughly $1 billion in annual free cash flow, for a stellar 31% margin.

With booming demand for ARM's data center CPUs, the company's revenue is expected to grow 23% in 2025 and 2026, up from 21% growth in 2023. This puts its 2026 expected revenue at $4.9 billion. These cutting-edge CPU chip architectures command above-average margins, which will propel ARM's profit margins toward a record high 48% in 2026. As a result, ARM's earnings per share are on track to reach $1.55 in 2024 and $2 in 2026, up from $0.32 per share in 2023.

Porter & Co.

PART 3: POWERING THE REVOLUTION

The Parallel Processing Revolution needs energy – yesterday. The Electric Power Research Institute estimates that data centers will consume roughly 9% of U.S. power generation by 2030 – more than double today's 4%. And parallel-processing data centers are a huge part of that surge... their combined energy consumption will be the equivalent of 30 million homes, about a quarter of all U.S. households.

Despite what climate change activists tell you, the only way to meet America's growing electricity needs will be to rely on traditional forms of cheap and reliable power, like natural gas and coal.

In the three reports in this section, we'll explore the powerful (and sometimes, dirty) forces needed to expand the Parallel Processing Revolution. Dark energy, indeed!

The first report in this section, **The King of Coal**, digs deep into one of America's oldest coal producers – and one of the most capital efficient. With its asset base of low-cost mines in the heart of America's richest coal basin, plus its strategically located Port of Baltimore marine terminal, **CONSOL Energy (CEIX)** has produced substantial free cash flows every year since going public in 2017. Plus, an upcoming merger in early 2025 will add another source of long-term growth from metallurgical coal – the kind used in steel-making. The merger will make the newly formed entity one of America's largest, lowest-cost producers of both thermal and metallurgical coal, securing decades of high-margin cash flows.

Aside from the parallel-processing boom, the company has a bright long-term future exporting coal overseas to developing nations. And better yet – fearful of political repercussions, most institutional investors won't touch this stock right now... meaning that it's trading at a deeply discounted valuation of just 5x free cash flow. In contrast, today's much-loved mega-cap technology stocks command valuations of 30x to 60x free cash flow multiples.

Our next report, **The Gods of Gas**, reveals a natural gas powerhouse that you've likely never heard of... **EQT (NYSE: EQT)**, a little-known company in Pennsylvania that's quietly becoming one of the largest suppliers of low-cost gas in America, and soon, the world. This leading "fracker" is a small, independent oil and gas firm whose production is centered on the largest natural gas reserve in the world, the Marcellus Shale. EQT sits on a resource that's so big... and is growing production so much... that it will become the world's most important energy company over the next decade. The company spent the last decade becoming America's largest independent gas producer, with a 17-fold increase in production since 2010. EQT will deliver natural gas directly from its wells through its pipelines to data centers in Virginia, making EQT's low-cost natural gas a crucial component of the Parallel Processing Revolution.

And, finally, **The Keystone** unveils the biggest power player of them all: nuclear energy. While the ruling class has no choice but to turn to fossil fuels for the next five to 10 years... they're desperately searching for an alternative energy source that's clean, cheap, and carbon-free. And there's only one viable candidate: nuclear.

BWX Technologies (NYSE: BWXT) has cemented itself as the key supplier of reactor design, components, and fuel for America's nuclear navy. Even more crucial, it's pioneering small modular nuclear reactors (SMRs), a powerful, portable energy source that could rip up the playbook of energy as we know it today.

Porter & Co.

THE BIG SECRET ON WALL STREET

The King of Coal

- Clean Energy Is A Scam
- Coal Rushes Into The "Green Vacuum"

FROM THE DESK OF PORTER STANSBERRY

The King of Coal

Generating the Energy for Today's Power-Hungry Data Centers

As the U.S. Abandons Coal, This Company Sends It Overseas

Judge Jacob Hessels spent much of the time in his courtroom dozing.

But he woke up when it really mattered… that is, whenever it was time to pronounce the death sentence.

Hessels couldn't really be blamed for taking a snooze while on duty. He was working round the clock, presiding over a seemingly endless lineup of heresy cases during the 1570s Spanish Inquisition in the Netherlands.

At the time, the Netherlands were under Spanish control… and Spain's Catholic ruler, Philip II, had given orders to round up all the Protestants in the Dutch "Low Country" and try them as heretics. But interrogating thousands of pious men and women – whose only crime was worshiping a little differently – proved exhausting for Judge Hessels.

After a few hundred cases, the judge started nodding off during the trials, waking up at the end to rubber-stamp executions with the words "Ad patibulum!" ("To the gallows with him!")

If a Protestant was lucky, he caught Hessels between power naps – and maybe managed to get his sentence commuted to lifelong banishment, with all his worldly belongings confiscated for the benefit of the Catholic Church.

Artist and engraver Theodor de Bry was one of the "lucky" Protestants spared the gallows.

De Bry forfeited all his possessions to the Inquisitors, and was exiled to Strasburg, Germany in 1570 with life and limb intact. The penniless – but talented -- De Bry got to work as soon as he landed in Germany… and before long, landed a history-altering art commission.

The project was a series of illustrations for travel books about the discovery and exploration of the New World (then, the destination of choice for gold-hungry Spanish conquistadors). De Bry's fanciful engravings included strange gods and sea monsters… and, on every page, depictions of Spanish adventurers chopping up and torturing innocent natives.

The illustrator's unfavorable – and realistic – portrayal of the Spanish had far-reaching consequences.

De Bry's books found an audience across Europe, and his vivid illustrations were reprinted and used in countless other books – including a widely-circulated tome by a conscience-stricken Spanish monk, Bartolomé de las Casas, who wanted to atone for the sins his countrymen had committed in the New World. In a way, De Bry's illustrations served as a kind of 16th-century Internet meme… a viral image that becomes a cultural touchstone.

During the same late 1500s/early 1600s time period, persecuted Protestants in several countries fought back against the Catholics…and eventually flipped the balance of power in Europe from Catholic to Protestant.

The kingdom of Spain, though, stayed stubbornly Catholic… and soon found out how it felt to be the underdog.

Protestant leaders like England's Elizabeth I and the Netherlands' William of Orange, along with powerful religious reformers like the French John Calvin and the German Martin Luther, launched a series of highly effective smear campaigns against still-Catholic Spain.

Fueled by recent memories of the Inquisition, and of course, by de Bry's lurid engravings, a popular belief took root: the idea that the Spaniards were cruel… dastardly… violent… and just a little bit worse than the rest of Europe.

That was the start of the "Black Legend" of Spain.

Black Legends and Black Diamonds

The Spanish "Black Legend" was bigger than any single propaganda campaign or smear tactic. Over many decades and across the continent – and world – it survived as a "big idea" that the country of Spain was, somehow, a global menace.

As Spanish historian Julián Juderías explained in his 1914 book *The Black Legend and the Historical Truth* (the first place we find the term "Black Legend" officially used), the Legend is "the systematic ignorance… of all that is favorable and beautiful in the various manifestations of culture and art, the accusations that in every era have been flung against Spain."

The Legend has taken many forms… subtle and not so subtle… over the last few centuries. (It's a poorly-hidden source for a lot of the woke "anti-Columbus" drivel that young people spout today.) At its roots, the Black Legend was a collusion – by powerful political groups and interests – to demonize a country that, while not perfect, really didn't deserve that level of vitriol.

Like all big stories, the Black Legend is a complicated one.

For one thing, there's more than a grain of truth in it. The Spaniards – while not the worst villains in European history by far – did torture their opponents during the Inquisition. They also undeniably committed genocide in the Americas – it's estimated that about eight million natives died, both from wartime violence and communicable diseases, during the three hundred years of Spanish colonization.

But as Juderías wrote… there's much that's "favorable and beautiful" in Spanish history, too.

The Black Legend is only one part of the complex story of Spain – and centuries of repetition have drowned out much of the other side… the good side… which has a legitimate claim to truth as well.

For instance, Spanish priests (fueled by the belief that they were doing the Lord's work) toiled tirelessly over the three hundred years of the Spanish Colonial period to bring Native Americans schools, churches, and hospitals. Spanish and Native populations frequently intermarried during that time. And – as a direct result of de las Casas' confessional writings – the Spanish government instituted a remarkable piece of humanitarian legislature, the Laws of Burgos, in 1512, ensuring that the conquistadors would treat the natives fairly, as free people, with cottages and land of their own.

In the end, there's a lot of gray in the Black Legend. And we would be unfair to Spain – and to history – if we accepted this massive defamation campaign at face value.

Right now, it's fair to say that we are in the midst of the creation of a new "Black Legend" for the modern age. We're seeing a loosely coordinated group of political powers determined to torpedo an imperfect – but still valuable, still useful – resource.

I'm talking, of course, about coal – which, appropriately enough, sometimes goes by the name "black diamonds." And about the "black legend" that's sprung up around it – not due to warring religious kingdoms, but to an even more fanatic group of zealots: climate change warriors.

That includes Barack Obama, who stated openly in his 2008 campaign that he wanted to "bankrupt" the coal industry – and followed through with a series of stringent actions to shut down coal mines and destroy jobs in the coal sector; President Joe Biden, who's continuing the war on coal with punitive EPA regulations designed to hamstring crucial coal mines in Montana's Powder River Basin; and environmental activist organizations like 350.org, who openly admit that they aim to conduct "smear campaigns" and to take away the industry's "social license to operate." All in favor of debunked green energy projects that are never going to work.

Coal is dirty energy, as we've written before. But it's also a vital part of American industry, and a significant segment of the country's energy backbone, responsible for 16% of power in the U.S., and 37% worldwide. And, especially now that we need massive doses of raw power to fuel AI's increasing energy demands… coal is not going away.

We owe it to ourselves not to believe the reductive myth-making of the climate-change elite. While nothing can whitewash the impact of coal pollution, we can't afford to ignore all the good things coal has done for the world, either.

In this issue, we'll show how one top-tier U.S. coal producer has thrived in recent years, despite the best efforts of American policymakers to bankrupt the industry.

Satisfying the Power Needs of Supercomputing

America's largest technology companies are competing in a capex arms race to overhaul their data centers for the new age of parallel processing.

This data center reboot comes as the explosion in artificial-intelligence ("AI") computing requires a shift from the previous regime of x86-based serial computing architectures to the new era of GPU-powered parallel processors – GPUs (graphics processing units) can carry out ten of thousands of operations at once with a processing capacity that is exponentially greater than previous chips. These high-powered GPUs will consume vast amounts of additional energy versus the traditional x86 processors of the past.

Before the parallel-computing revolution took hold, the typical rack of servers in a data center required about three to five kilowatts (kw) of power. These devices run 24 hours a day, consuming 72 to 120 kilowatt-hours of electricity each day, or about three times as much as an average U.S. household. In today's cutting-edge data centers that run on high-powered Nvidia GPUs, this number has increased 10-fold to roughly 30 kw to 50 kw per rack. As a result, the GPU-based data centers of today's parallel-computing revolution consume enough electricity to power 50,000 homes.

The Magnificent Four mega-cap U.S. tech companies – Alphabet (GOOG), Amazon (AMZN), Meta (META), and Microsoft (MSFT) – are racing to overhaul their roughly 600 U.S. data centers with the latest energy-hungry Nvidia GPUs. Together, these companies will spend a record $168 billion on capital expenditures this year, up 50% from 2023, with the majority going toward data center investments. And that's only the beginning of a $1 trillion investment boom that analysts expect U.S. companies will direct toward new data center development over the next five years.

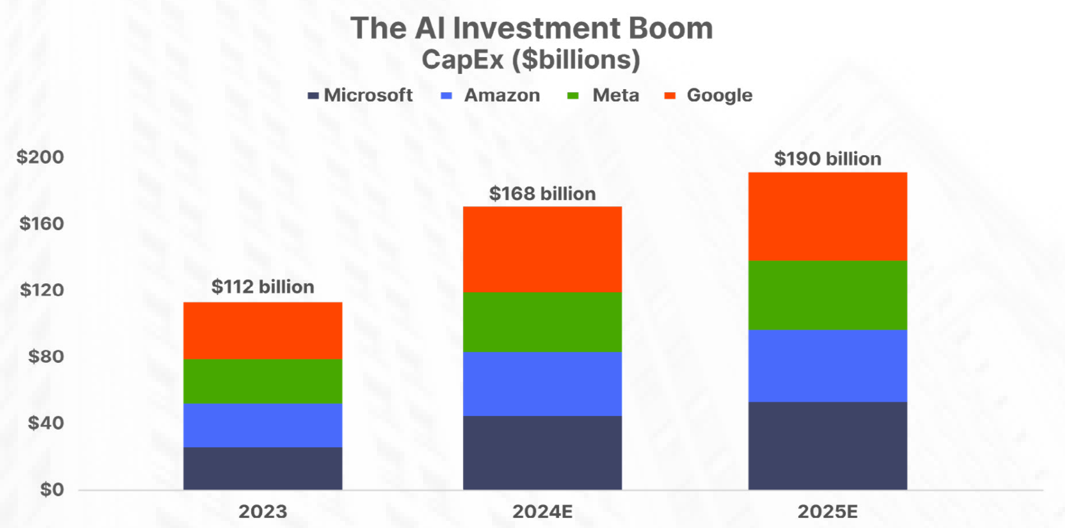

Source: Bloomberg

The Electric Power Research Institute estimates that data centers will consume roughly 9% of U.S. power generation by 2030 – more than double today's 4%. Combined with the rising power demand from electric vehicles, this surge in demand will strain America's aging power grid. Indeed, it's already happening.

U.S. policy-makers plan to fulfill America's growing electricity needs with "renewable" power sources like wind and solar. The federal government currently aims to make 80% of America's electric grid powered by renewable sources by 2030, up from just over 20% currently. This would require solar and wind power to reach a combined output of 6.8 trillion kwh of annual generation capacity – or a 10-fold increase over current levels – by 2029.. Even assuming the manpower and materials existed to make this happen, it would require an investment on the scale of many trillions of dollars.

Even die-hard environmental activists, like Ernest Moniz, the secretary of energy in the Obama administration, acknowledge the futility of relying on wind and solar for America's growing power demands. At a **power-industry conference in March**, Moniz explained:

> *"We're not going to build 100 gigawatts of new renewables in a few years."*

Moniz further explained that the only way to meet America's growing electricity needs will be to rely on traditional forms of cheap and reliable power: most notably, natural gas and coal.

The challenge is that U.S. utility operators haven't upgraded America's powergeneration capacity for the last 15 years, as electricity demand has been stagnant over that time. As a result, the hundreds of billions of dollars pouring into new data center construction is rapidly outpacing the growth in new power generation capacity. Commercial real estate firm CBRE recently noted that the lack of power supply has delayed construction of data centers by two to six years.

This all adds up to a looming energy crunch. But Bloomberg reports because of increased demand, the push to shutdown coal plants is slowing down. In fact, the closing of more than 20 facilities from Kentucky to North Dakota that were set to retire between 2022 and 2028 has been delayed.

In this issue, we'll show how one U.S. coal producer is poised to profit from supplying the critical power fueling this computing revolution.

A Highly Overlooked Parallel-Computing Play

Headquartered in Cecil County, Pennsylvania – just outside of Pittsburgh – **CONSOL Energy (NYSE: CEIX)** is one of America's oldest coal producers. Its parent company Consol first began mining in 1864 in the Appalachian Basin, one of America's richest coal deposits. CONSOL Energy was spun off and began trading as a public entity in 2017.

On August 21, 2024, thermal coal producer CONSOL Energy (CEIX) announced a merger agreement with fellow coal miner Arch Resources (ARCH) to form a new entity called **Core Natural Resources (NYSE: CNR)**. The deal terms state that CONSOL shareholders will own 55% of the merged entity, with Arch shareholders owning the remaining 45%. **Please note, the deal closed on January 15, 2025, and CONSOL shares were converted to the new entity Core Natural Resources.**

CONSOL Energy primarily produces thermal coal, which is used to generate electricity. In 2023, the company mined 23 million tons of thermal coal, making up 88% of its total output. The remaining 12% of production was metallurgical coal, which is used in steel-making.

In total, the company mined 26.1 million tons of coal in 2023, up 9% from 2022. This brought in $2.57 billion in revenue and $687 million in free cash flow, for a stellar free cash flow margin of 27%. For perspective, that's on par with some of the world's most dominant technology giants, including Microsoft, Meta, Apple, and Alphabet (Google):

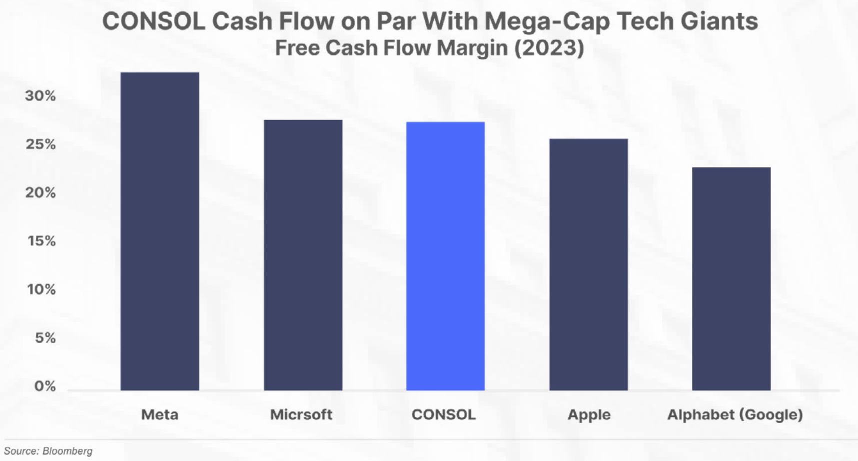

Despite CONSOL's impressive profit margins, it gets little respect on Wall Street. Many hedge funds, pension funds, and market index providers simply won't own coal stocks for fear of being labeled "non-woke."

At Porter & Co. we're more than happy to capitalize on the folly of the woke ideologues. Wall Street's dislike of coal stocks has created a tremendous opportunity in shares of CEIX, which trade at a deeply discounted valuation of just 5x free cash flow. In contrast, today's much-loved mega-cap technology stocks command valuations of 30x to 60x free cash flow multiples. While there are admittedly big differences between coal-mining and technology companies that warrant different industry valuations, the cash these industries generate still spends the same.

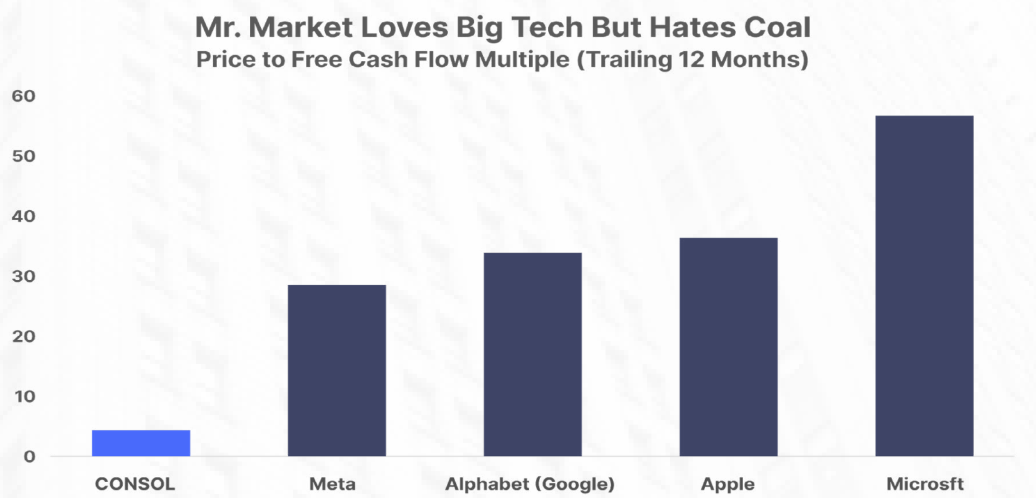

Source: Bloomberg

This dirt-cheap valuation makes CONSOL one of the most unloved, underappreciated winners of the parallel-computing revolution.

The key thing to understand about investing in U.S. thermal coal producers is that domestic consumption has been on a steady downward trend for nearly two decades. This has given rise to a widespread misperception that U.S. mining is a dying industry that will soon go extinct.

But the truth is, this industry backdrop has enabled top tier producers like CONSOL to thrive. To understand why, let's briefly review the supply/demand trends in the overall market.

Coal Is Dead... Long Live the Coal Miners

Unlike metallurgical coal, which has no viable substitute for its use in steelmaking, there are a number of alternatives to thermal coal – natural gas, nuclear power, wind generation, and solar – all of which produce electricity with significantly fewer (or zero) direct carbon emissions.

Since 2008, environmental mandates have forced U.S. utilities to shut down hundreds of coal plants in favor of these lower-carbon alternatives. As a result, U.S. thermal coal consumption used in power generation has declined by roughly 60% since 2008, resulting in a similar decline in mining output:

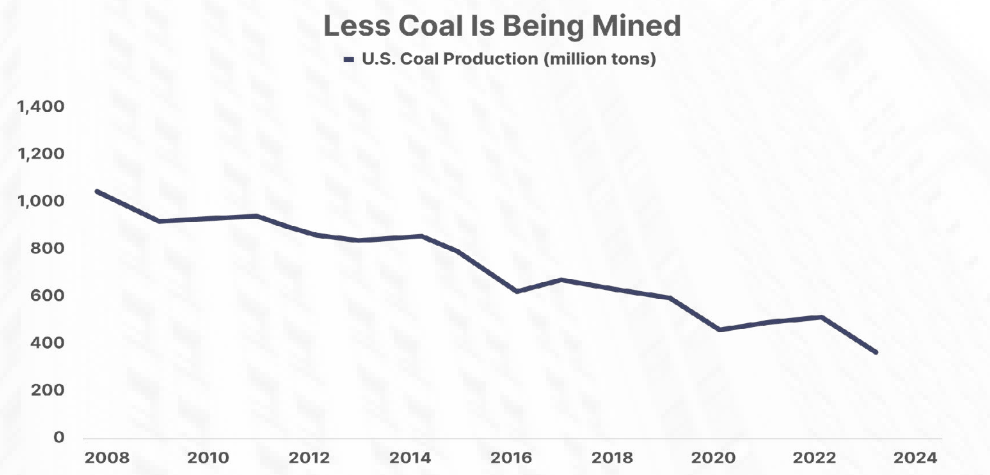

Source: Energy Information Administration (EIA)

CONSOL has remained immune from the broader decline in the U.S. domestic coal market by tapping into overseas markets, with exports making up nearly two-thirds of its sales (discussed in greater detail below).

At the same time, there's a silver lining within this industry decline. For the best-inclass U.S. coal miners, declining domestic demand has eliminated the competition and the excess supply coming from the weaker players in the industry. The chart below shows how the number of active U.S. coal mines has fallen by 72% since 2008:

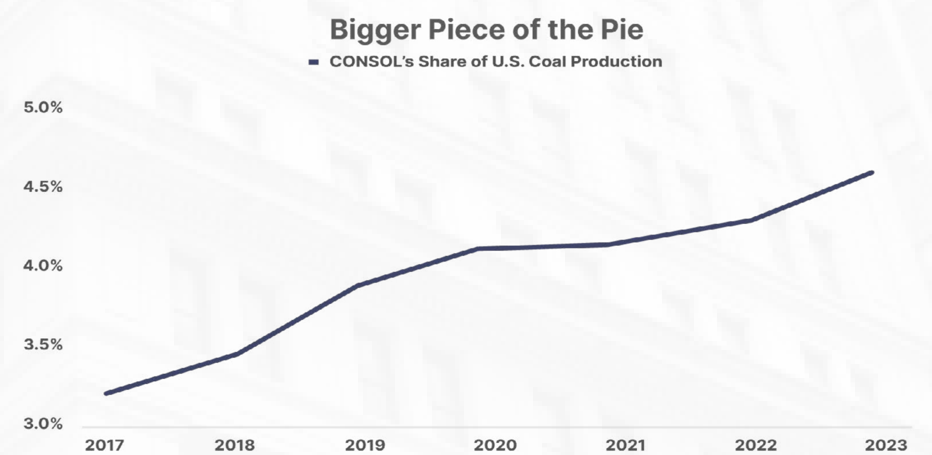

Source: Energy Information Administration (EIA)

One of the key factors behind CONSOL's resilience is the superior quality of its mining assets.

The Secret to CONSOL's Success: Energy Density

CONSOL's core cash cow, responsible for 99% of its coal production, is a collection of mining properties that make up the Pennsylvania Mining Complex ("PAMC"). The PAMC assets include the Bailey, Enlow Fork, and Harvey mines, each located in the southwestern corner of Pennsylvania. Together, these three mines generated a total of 26.1 million tons of coal last year:

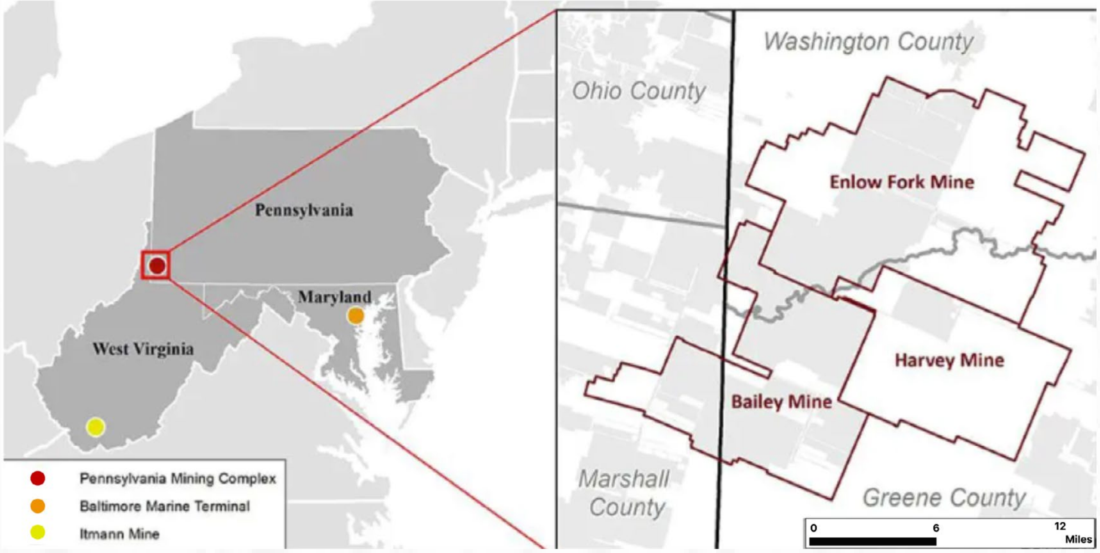

The first key feature of these three mines are their location, in the northern part of the Appalachian coal basin. Appalachia is home to some of the most energy-rich thermal coal in the U.S., and the energy density increases toward the northern end of the basin.

CONSOL's PAMC mines in northern Appalachia contain some of the most energy-rich coal among all major U.S. coal basins, with an average of 12,972 British thermal units (Btu) per pound. This is up to 50% more energy-rich than other major U.S. thermal coal basins, including Central Appalachia (in Kentucky, Virginia, West Virginia, and Tennessee), the Illinois Basin (in Illinois, Indiana, and Kentucky) and the Powder River Basin (in Montana and Wyoming):

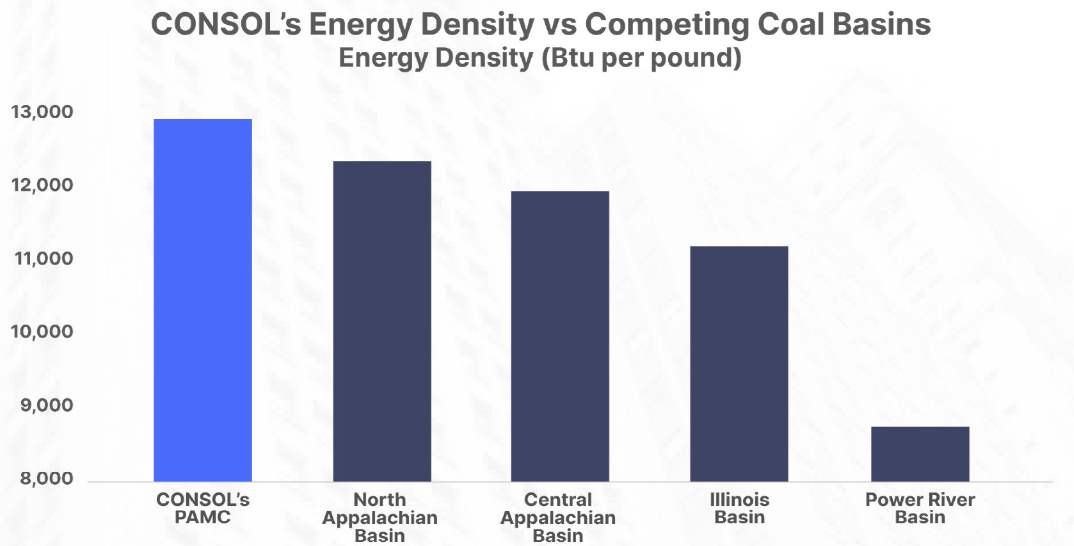

Source: Consol Fillings

Energy density is critical for two reasons. First, energy-rich coal commands premium pricing because it generates more electricity per pound. The second advantage is that denser coal produces more energy per unit of carbon dioxide emissions – a byproduct of coal combustion that utilities seek to minimize, in order to comply with environmental regulations. CONSOL's energy-rich coal allows it to sell to utility operators that operate some of the cleaner-burning coal plants. This provides it with a more stable customer base that's at lower risk of going out of business compared with its competitors that sell less energy-rich coal to utilities with higher emission profiles.

Six of CONSOL's top domestic power plant consumers, which each purchase over 500,000 tons of coal each year (2% of CONSOL's output), have been customers for at least five consecutive years. Securing long-term relationships with stable utility operators is one key reason why CONSOL has managed to grow its coal volumes more than 5% from 24.8 million tons in 2017 to 2023, even as the overall industry has suffered double-digit volume declines over that same period.

A Low-Cost Leader

The other advantage of CONSOL's PAMC mines is that pulling coal from this area is a low-cost operation. That's because of the Pittsburgh Number 8 Coal Seam, the geological formation that hosts PAMC's energy-rich coal. Coal seams are long, continuous rock formations that enable what's known as "longwall" mining. This process involves shearing off large, slices of coal-containing rock from a long face (or wall) of a geological formation in a single pass.

Compare this with underground mining methods, like the "room and pillar" method, which requires digging into an underground coal formation, creating a mining "room" for extracting coal, while leaving large "pillars" of the formation in place to support the roof.

Longwall mining offers a safer, faster and more economically efficient approach than alternatives like room-and-pillar mining. By eliminating the need to dig elaborate underground structures, it requires less man and machine power per unit of coal mined. Longwall mining also extracts a higher proportion of the available coal in the formation versus underground mining, which leaves substantial chunks of the coal in place to prevent the structure from collapsing. As a result, underground mining methods like room and pillar typically only recover about 50% to 60% of the available coal deposit, compared with more than 75% that longwall mining extracts.

Finally, with a history of mining Appalchian coal since 1864, CONSOL has 160 years of experience, helping it optimize its production methods. All of these factors together contribute to CONSOL's industry-leading profitability.

In 2023, the cost of producing coal from CONSOL's PAMC assets averaged $36.10 per ton, versus a selling price of $77.74 per ton. With a $41.64 cash profit per ton, CONSOL's PAMC assets delivered an operating margin of 54% – making it one of the most profitable commodity producers in the world.

The reserve base for CONSOL's PAMC assets ensures a long future of profitable production. As of year end 2023, the company had 583.5 million tons of reserves among its three PAMC mines. Pulling out its current level of 26 million tons a year leaves enough to support 22 years of future production.

And CONSOL's future as a coal miner isn't limited to a declining U.S. market. Coal offers the cheapest, most reliable form of baseload energy in the world. So even as climate alarmists have hurt demand in the U.S., the rest of the world continues building new coal-fired power plants. That's why coal consumption recently hit all-time highs in 2022, following a brief decline in 2020-2021 when global economic activity slowed during the COVID-19 pandemic:

Coal Is the Super Power Generator

Global electricity production by source, in terawatt hours (twh)

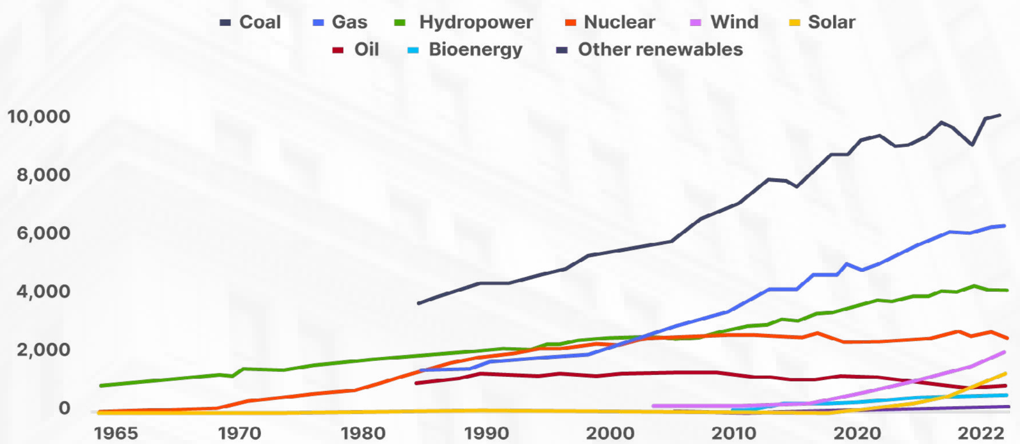

Source: Ember, Energy Institute

Most experts forecast that global coal demand for power generation will continue growing through at least 2030. The key drivers of this demand will come from India and China, which together make up one-third of the world's population. In 2023 alone, China grew its coal capacity by 48.4 GWs (or 4%), while India added 14 GWs of capacity.

The massive (and growing) global market for thermal coal means that, even in the ultimate bear case scenario where U.S. coal consumption goes to zero, CONSOL can still thrive.

That's all thanks to CONSOL's ownership of a key export terminal, strategically located to deliver a low-cost advantage for coal shipments into international markets.

How CONSOL Gets Coal to the World

Perhaps CONSOL's most valuable asset is its Marine Terminal, located in the Port of Baltimore. This is the only major east coast coal terminal that can receive shipments from two railroad operators, Norfolk Southern and CSX. The terminal includes 19.3 miles of railway track, with three railway sidings, which are short tracks used for loading and unloading freight that do not interfere with the main line operations. This infrastructure allows the CONSOL Marine Terminal to seamlessly load coal directly from rail cars to shipping vessels, with a capacity of up to 9,000 tons of coal per hour (or 78 million tons per year, or around three times CONSOL's total annual production).

Ownership of this port infrastructure provides several key competitive advantages for CONSOL. The first advantage is the port's location, only 250 miles from the PAMC mining complex in Northern Appalachia. This short hauling distance provides CONSOL with one of the lowest transportation costs for moving its coal from the mine site to export ships, at an average cost of just $18 to $19 per ton, which is 15% to 20% cheaper than the $21 to $25 cost of moving coal from other key U.S. basins to alternative export terminals:

The Cost to Move Coal

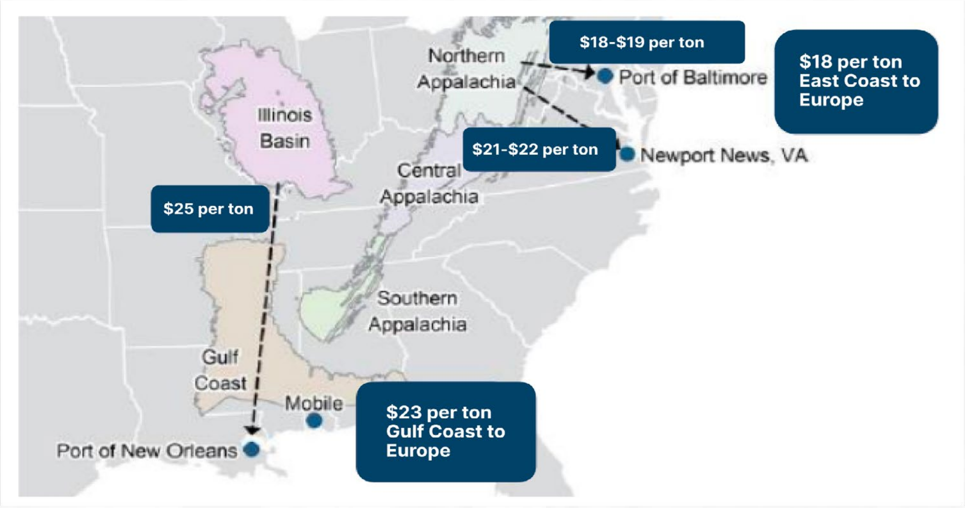

The second key advantage is that CONSOL doesn't have to pay a third party to handle its coal at the terminal. Instead, the Marine Terminal has enough capacity to load coal from other mines onto shipping vessels, earning a substantial stream of revenue. In 2023 CONSOL earned $106 million in terminal revenue, plus $294 million in freight revenue from the Marine Terminal. Together, these two revenue streams generated $400 million in sales, or 16% of CONSOL's total $2.6 billion in 2023 revenue.

What's more, this terminal provides CONSOL with a low-cost advantage for routing its domestically produced coal abroad. Since 2017, CONSOL has doubled its export coal sales from 8.3 million tons (or 32% of its total volume) to 16.2 million tons in 2023 (or 61% of total volumes).

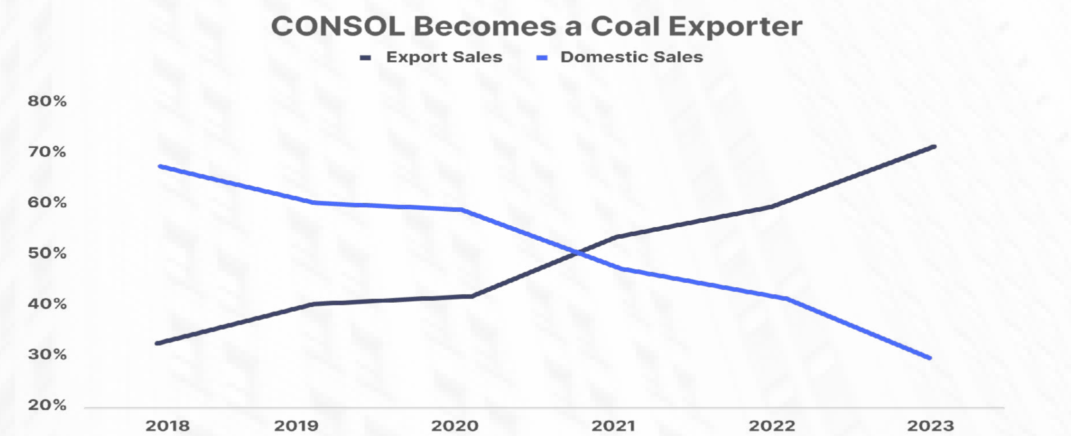

Source: Consol Fillings

And therein lies the Big Secret about CONSOL: the company is rapidly changing itself from a domestic U.S. coal supplier into a coal exporter. And while it will enjoy a boom in domestic demand for the data center build out, its long-term future is secured as one of America's leading, low-cost exporters of high-quality coal to a large and growing international market.

A Diamond in the Rough

With stock prices trading at record-high valuations across the board, CONSOL is an exception, offering a rare, deep value opportunity.

Its current share price of around $100 gives CONSOL a market capitalization of $3 billion. Its quarterly results are volatile given the swings in coal prices in recent years, amplified by the Russian invasion of Ukraine that created a short-term coal shortage and price spike in 2022. However, at current coal prices the company should be able to generate roughly $450 million to $500 million in free cash flow each year.

With its rich asset base of low-cost mines in the heart of America's richest coal basin, plus its strategically located Port of Baltimore marine terminal, CONSOL has produced positive free cash flows every year since going public in 2017 – including in 2020, when coal prices collapsed during the COVID-19 outbreak.

The company's steady cash flows support a pristine balance sheet with more cash than debt, allowing it to funnel its excess cash to investors. CONSOL's shareholder-friendly management team has committed to returning 75% of cash flow to investors, primarily in the form of share repurchases. Just in two years, 2022 and 2023, the company has reduced its share count by nearly 20%.

The combination of growing production volumes and cash flows, plus a falling share count, has propelled CONSOL's earnings per share nearly 10-fold since going public, from $2 in 2017 to $19.64 in 2023.

Even though the company gets little attention on Wall Street, and thus trades at a deeply discounted valuation, its share price has handily outperformed the market. Since 2017, CONSOL has delivered a 364% total return, or 27.6% compounded annual growth rate ("CAGR"). Over the same period, the S&P 500 has gained 106%, or a 13.6% CAGR:

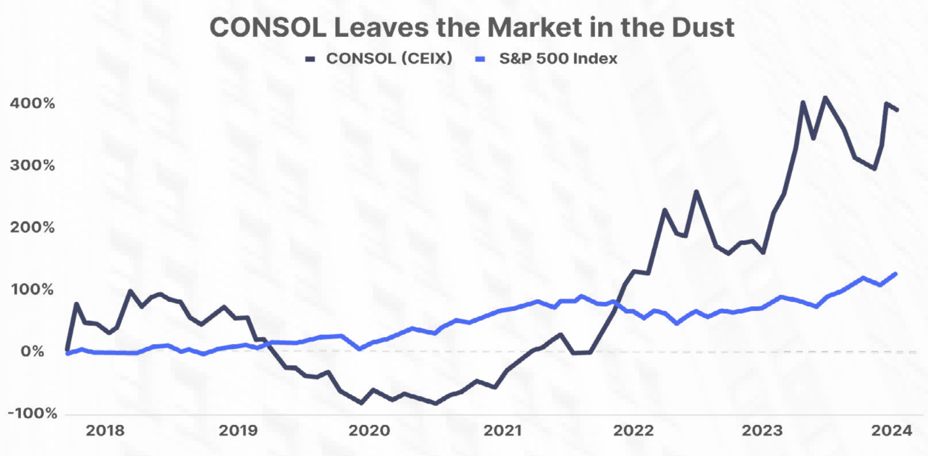

Source: Bloomberg

Given its deeply depressed valuation of just under 5× 2023's free cash flow, CONSOL is our number-one pick to capitalize on today's parallel computing boom – and the only company in this segment of the Big Secret portfolio that we recommend buying at current prices.

CONSOL typically trades in a range of 4x to 8x cash flows, with an average multiple of 6x since its inception as a public company in 2017. However, this turns out to be a boon to investors, since it allows the company to repurchase a greater number of its shares versus companies that trade at higher valuations. Over time, the combination of growing earnings and cash flows over a rapidly shrinking share count should provide a wonderful formula for shareholder returns.

Porter & Co.

THE BIG SECRET ON WALL STREET

The Gods of Gas

- You've Never Heard of These Two Brothers...
- But They're About to Transform the Global Energy Market

FROM THE DESK OF PORTER STANSBERRY

The Gods of Gas

You've Never Heard of These Two Brothers...

But They're About to Transform the Global Energy Market

The company's corporate slogan should be "death to Saudi Arabia."

While the media and our political leaders are fixated on the impossible: a world without fossil fuels...

...Two brothers from Pittsburgh pieced together the world's leading producer of an "ESG-approved," super-clean, carbon-based energy source.

Now they're building a dominant distribution network that will render Saudi oil virtually worthless.

The biggest disruption to the world's energy markets in 100 years is underway. And you've never heard about it.

Until now.

"Bring me the head of the dog," said the angry voice on the speaker phone.

The voice was the chief of staff to Mohammed bin Salman, the acting king of Saudi Arabia. The phone was inside a safe room at Istanbul's Saudi Consulate. The order wasn't rhetorical: the leader of Tiger Squad, an elite 15-member hit team, was holding a surgical bone saw.

Jamal Khashoggi, a *Washington Post* journalist, was about to die. While his girlfriend waited outside the consulate, the Tiger Squad put a plastic bag over his head and cut him into pieces.

According to the recording of his death, Khashoggi took seven minutes to die, while Mohammed bin Salman listened. The Tiger Squad brought back Khashoggi's fingers, which they cut off one at a time as a trophy for the king.

The rest of Khashoggi's body was burned in a specialized consulate oven. A courtyard barbecue was used to mask the smell.

The depravity of Saudi Arabia's leaders isn't new.

Barbaric events occur routinely in the country. People convicted of adultery are stoned to death. Amina bin Salem Nasser was beheaded for practicing sorcery and witchcraft in 2011. (We doubt she was really a witch. Guess we'll never know for sure.)

And let's not forget the mass executions. Last March, Saudi Arabia simultaneously beheaded 81 people. They were all confessed criminals... because they had all been tortured.

There's only one difference between Saudi Arabia and the barbaric, friendless countries of the Middle East, like the Taliban's Afghanistan—oil.

Since the end of World War II, Saudi oil has powered much of the world, with tankers carrying its crude oil to virtually every corner of the earth.

For the last 50 years, the quest for energy security has dominated foreign policy of all the great powers, including the United States. This has meant that Saudi Arabia has always been given a pass. Even the country's obvious ties to the 9/11 terrorists were swept under the rug.

But the world's market for energy is about to be turned upside down—forever. And in another decade, no one will need Saudi oil anymore.

The Big Secret on Wall Street this week isn't about America's reliance on imported petroleum products to keep our economy running. Everyone already knows there's not enough refining or pipeline capacity in America. We know about the Jones Act (which restricts the shipping of petroleum products inside the U.S.).

So... why haven't more refineries been built? Why haven't more oil pipelines been laid? Why don't we repeal the Jones Act?

Because the gasoline business is going extinct. It's a dinosaur. Automobile manufacturers are standardizing on electric cars. Thus, capital invested in new gasoline refining and distribution will end up being "stranded"—unused, forgotten, and worthless. That's why big investors won't get behind any new gasoline infrastructure.

Gasoline is NOT the gas you should focus on.

The Big Secret on Wall Street this week isn't about gasoline – it's about gas. Natural gas. And there's one little-known company in Pennsylvania that's quietly becoming one of the largest suppliers of low-cost gas in America, and soon, the world.

This leading "fracker" is a small, independent oil and gas firm whose production is centered on the largest natural gas reserve in the world, the Marcellus Shale.

Through a series of strategic acquisitions, this company is transforming into the first super-major energy company to emerge from America's shale resources, which are the largest ever discovered.

This firm (which we'd bet not a single paid subscriber has previously heard of) has suddenly – virtually overnight – become the largest producer of U.S. natural gas. It will soon be the world's largest and most important energy company. Read that again.

A company you've never heard of before is already the leading producer of natural gas in the United States.

This company sits on a resource that's so big and is growing production so much that it will become the world's most important energy company over the next decade. The company spent the last decade becoming America's largest independent gas producer, with a 17-fold increase in production since 2010:

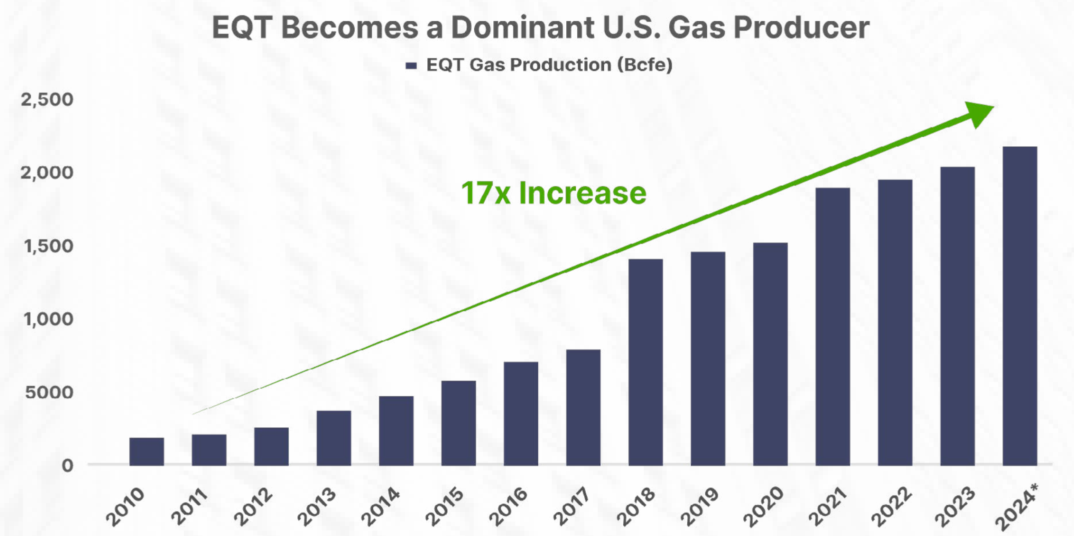

Source: Company Filings; Management Guidance for 2022

Best of all, the revolution that this company is leading will render Saudi oil much less important than it's been in the past. Its corporate slogan ought to be "death to Saudi Arabia."

Meet The Rice Brothers — "The Gods Of Gas"

Toby and Derek Rice are from Pittsburgh.

Their father was a private equity banker who specialized in oil and gas. The brothers, while still in their late 20s, began to assemble valuable acreage in the Marcellus shale basin, starting in 2007. When the emerging shale field produced surplus amounts of gas, they took advantage of collapsing prices to add huge amounts of acreage from failing producers. Their privately-owned firm, Rice Energy, grew to be one of the ten largest natural gas producers in the U.S.

But that was just the beginning.

In 2017, the brothers sold Rice Energy for $6.7 billion to a large, publicly traded gas company, EQT **(NYSE: EQT, $50)**, creating the largest producer of U.S. natural gas.

Then, in 2019, unhappy with EQT's inability to control costs or increase production, the brothers conducted a proxy battle and won 80% of the votes. Toby became EQT's CEO, which is when the story gets interesting.

By the end of 2021, despite the Covid-19 disruptions, the Rice brothers had grown EQT into one of the world's most efficient energy companies—including changes that were almost too good to be true. In just over a year, well costs fell 47%, and drilling speeds increased by 95%. The financial impact was substantial.

Gross profit margins more than tripled, from less than 10% to over 30%.

Cash from operations grew from $1.5 billion to over $3.5 billion, even though gas prices remained low and capital expenditures were flat.

Earnings per share went from negative $.89 in 2019 to positive $.89 in 2021.

Since then, the earnings have increased to 2.28 per share in 2023.

Most importantly, free cash flow—the excess capital available to return to shareholders—exploded.

In the five years from 2015 - 2019, under the old regime, EQT reported total free cash flow of negative $3.4 billion. In the last five years from 2020 - 2024, under the leadership of the Rice brothers, free cash flow surged to $5.0 billion.

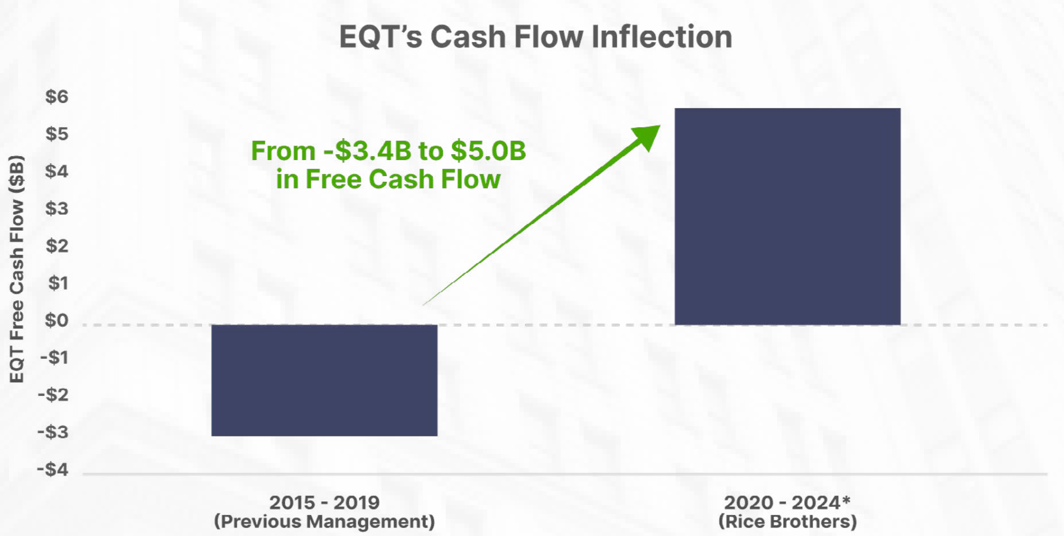

Source: Company Filings; Management Guidance for 2024

Over the next five years from 2025 - 2029, the company expects to generate between $8 billion and $27 billion in cumulative free cash flow, based on gas prices ranging from $2.75 - $5 per million British thermal units (MMbtu). For perspective, that's roughly 40% - 135% of its current market capitalization of just $20 billion.

How's that possible? How did two brothers from Pittsburgh take a small, regional, "also-ran" shale gas company and turn it into an economic engine that produces tens of billions of free cash flow and is the largest producer of natural gas in the U.S.?

The brothers didn't merely cut costs—they also struck a great deal with Chevron. And that deal was huge.

In the fall of 2020, the pandemic sent oil and gas prices to decade lows. Chevron wrote off its entire Marcellus investment. In fact, Chevron took an $8 billion write-off.

But EQT paid only $735 million for Chevron's Marcellus operations. They practically stole it.

The Chevron deal (800,000 acres) and EQT's later acquisition of Alta Resources (another 300,000 acres) assures that EQT will remain the dominant provider of Marcellus natural gas for decades.

Estimating what that will mean over time is difficult, thanks to EQT's continuous drive to improve operating efficiencies. It's also hard to know just how much gas EQT controls. But a safe bet is a lot more than the 27 trillion cubic feet that have been proven with current drilling.

As drilling techniques improve and more wells are drilled, the size of the total Marcellus resource continues to scale higher and higher. As of 2019, the United States Geological Survey (USGS) estimated that the Marcellus formation (including the associated shale layers known as the Utica) contains 214 trillion cubic feet of natural gas. However, these same estimates have been increasing over time, from 2 trillion 20 years ago to 84 trillion 10 years ago to 214 trillion most recently.

Scientists from Penn State University now claim that over 400 trillion cubic feet of gas is recoverable in the basin. But the truth is, no one really knows for sure.

To put this into context, the Marcellus probably contains more natural gas than all the other natural gas producing areas in the U.S., combined. The Marcellus, alone, probably contains more natural gas than every other producing nation except Russia, Iran, and Qatar.

The Marcellus isn't merely a big gas field. It's one of the largest reservoirs of energy in the world. Its development will not only change the U.S. economy, but it will also reshape the global energy map for the rest of our lives. And EQT will lead this process – because no one will produce more natural gas from the Marcellus (or in America) than EQT.

And that's not just because EQT owns the biggest acreage position in the heart of America's most prolific gas basin. It's because we can count on the Rice Brothers to continue expanding that footprint through savvy deal-making.

In September 2022, the Gods of Gas stuck again with another key strategic acquisition – the $5.2 billion purchase of oil and gas assets of Tug Hill, a midsize gas producer in the heart of the Marcellus. The acquisition expanded EQT's Marcellus footprint by roughly 10% to 1.1 million acres, and boosted its average daily production to 6.3 billion cubic feet equivalent – a 15% increase.

But perhaps most important is the quality of this acquired acreage, which shows up in the rock bottom breakeven costs of just $1.35 per million British thermal units (MMBtu). This reduced EQT's breakeven costs from $2.30/MMBtu to $2.15/MMBtu.

As part of the deal, EQT also acquired XcL Midstream, a natural gas pipeline company with an extensive network throughout the Appalachian Basin which is home of the Marcellus. With these new pipeline assets, EQT will expand the scale of its operations towards becoming one of the world's largest, vertically-integrated natural gas supermajors.

On another front, a critical first step toward becoming a super-major energy company is gaining an investment grade credit rating. This will allow EQT access to much more capital, which it will need to build out more pipelines, more processing plants, and, eventually, its own LNG infrastructure (liquified natural gas – which we'll talk about more below).

EQT received an investment grade credit rating from both S&P and Fitch in 2022. In August 2023, Moody's followed suit and raised EQT's credit rating to investment grade as well.

EQT was essentially forgotten and left for dead during the pandemic. But today, it has the scale, market power, and credit rating to do something only super-major oil companies can do—build its own global distribution network and capture the vastly higher prices for energy on the global market.

Over the next decade, EQT's best-in-class natural gas acreage, pipelines, processing plants, and long-term, fixed-priced global distribution deals will become the envy of every energy company in the world. But until you read this report, you'd never even heard of EQT or the Rice brothers—right?

That's because the media and politicians are, as always, fighting the "last war." They play to the plebes who care about filling up a SUV. Think about all the Biden stickers on gas pumps: Those stickers are there because the media and politics focus on today's problems.

But the future is obvious. Gasoline isn't going to power the world's transportation economy for the next 50 years. Global automakers are investing hundreds of billions in vehicle electrification over the next decade.

In 2022, more than 10 million EVs were sold globally – making up 14% of all new car sales. The International Energy Agency (IEA) expects this number to reach 17 million in 2024. And by 2035, the IEA expects more than half of all new car sales will be electric.

That's why nobody wants to own a new gasoline refinery (with a 30-year useful economic life). Demand for gasoline is going to fall off a cliff in less than a decade.

The next gasoline-powered car you buy will very likely be the last gasoline-powered car you will ever own.

As electricity replaces gasoline in vehicles, the ultimate fuel source for cars will change from gasoline to natural gas. Natural gas will power the electric grid, not gasoline. If you want to plug your car in, you're going to need what EQT has – and lots of it.

> *What investors need to know isn't what the price of gasoline is going to do by the end of this year. What you need to know is how America's dominance in natural gas is going to completely reshape the market for energy and transportation all over the world.*

If you followed our work at Stansberry Research, you know we've been covering the shale revolution for over a decade.

You also know that we broke some of the biggest stories in finance for years, such as predicting the collapse of Fannie Mae and Freddie Mac, GM's bankruptcy, and the demise of GE.

We have also recommended dozens of great emerging companies that went on to become industry leaders, such as Amazon, Qualcomm, Illumina, Microsoft, Shopify, Nvidia, and literally dozens more.

But what's about to happen with U.S. natural gas is far bigger than any of these things.

American natural gas is emerging, right now, as the world's next dominant energy source.

Forget about Saudi Arabia. America is the new energy king. And there's one company best positioned to capture the biggest profits of this new global reality: EQT.

A new super-major energy company is emerging—the first all-American corporation that can frack, refine and distribute low-cost natural gas from the world's largest natural gas field (the Marcellus) to virtually any country in the world.

EQT: The King Of Natural Gas

What's the richest country in the world on a per capita basis?

Lots of people would guess Saudi Arabia. Or maybe Kuwait. Or the United Arab Emirates. But it's none of those countries – it's Qatar.

Qatar was a relatively poor country until the early 2000s, with a GDP below $10 billion.

However, beginning in 1997, Qatar quietly came to dominate the world's global trade in LNG. Qatar shares a huge offshore natural gas field with Iran, known as the North Field.

The field is an enormous resource—one of the world's largest proven natural gas fields, with reserves of at least 896 trillion cubic feet (tcf). But Qatar didn't begin exporting natural gas in large quantities until 1997, sending its first LNG shipment to Spain.

By 2007, Qatar was the world's largest LNG supplier. Today, Qatar has eight massive LNG "trains" and six even larger "mega-trains," which can liquify huge volumes of natural gas for shipment on specialized LNG tankers. Qatar is currently investing another $30 billion in a massive North Field expansion, which will reportedly increase production by 40% by 2025.

The results of these investments are hard to believe.

Qatar's GDP grew from $9 billion annually in 1996 to over $200 billion in 2014. Qatar's economy grew 21-fold in less than 20 years. The nation's sovereign wealth fund now tops $500 billion in assets under management, making it one of the world's largest capital pools. With only 300,000 citizens, Qatar has a per capita GDP of $686,000, and more than $1 million for each citizen in its sovereign wealth fund.

That's the kind of wealth that's coming to America.

How do we know?

The U.S. began exporting significant natural gas quantities in the early 2000s via pipelines to Canada and Mexico. As U.S. production grew thanks to shale gas development (resulting in the U.S. becoming the world's largest natural gas producer in 2009), exports increased rapidly. Exports grew from less than half a billion cubic feet daily in the early 2000s to over two billion cubic feet daily in 2015.

Since 2015, export growth has been parabolic – growing more than 4-fold from less than two trillion cubic feet to nearly 8 trillion cubic feet per year.

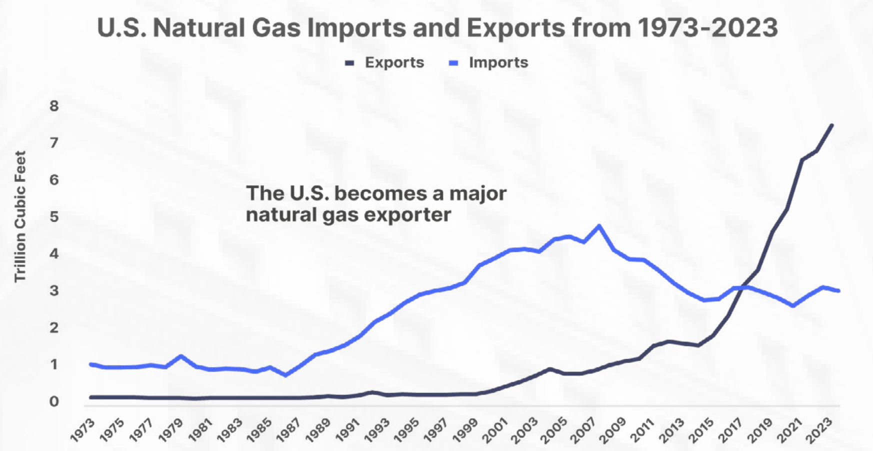

Source: Energy Information Agency

Longtime readers of our work may remember a report we wrote in the spring of 2006, titled "Madness."

The report was about a start-up that planned to build a huge new LNG import terminal in Louisiana. This was during the "peak oil" mania, when most investors sincerely believed the U.S. was running out of hydrocarbons (oil and gas) and would face permanent shortages.

Some argued the "only way" to save the country was by importing huge quantities of oil and gas from places like Russia and Qatar, where major oil companies were investing tens of billions. Some of these projects were incredibly risky—even stupid. One involved natural gas production in the middle of the Caspian Sea. Another project was in the Russian arctic, 300 miles from the North Pole!

It was a global mass hysteria. And frankly, we couldn't understand why everyone had lost their minds. We knew America had more hydrocarbons locked in so-called "tight shales" than these other places combined. All we needed were some pipelines and a little ingenuity.

We believed, even back then, that America would be the dominant provider of natural gas to the world—not an importer. As we saw shale gas drilling begin to take off, we also saw more and more gas being produced and stored. A glut was forming, not permanent shortages.

As we wrote back in May 2006:

> *"I believe over-investment in domestic drilling and production has already produced a glut of natural gas that will persist for many years... New technologies recently have unlocked huge supplies of gas in the United States. Heavy investment in the sector since 2003 is now beginning to bring these new reserves into production. Far from running out of natural gas, we're drowning in the stuff. Huge new supplies of gas have been found in the U.S. over the last 10 years because of innovations in shale-gas drilling. These new reserves are only now coming into production...*
>
> *"[Cheniere Energy plans] to build three new liquified natural gas facilities along the Gulf coast. Each of these terminals will cost more than a billion dollars. They will take several years to construct. The first one isn't scheduled to begin operations until 2008 or 2009. The point of these terminals is to serve as off-loading stations for LNG tankers, which, theoretically, would ship natural gas to America from places like Egypt, Algeria, and Oman. Cheniere wants to spend billions to set up facilities for importing natural gas into the United States. This is utter madness. There's only one other country in the world, according to the C.I.A., that produces more natural gas than the United States: Russia. Cheniere's business plan is the equivalent of setting up a really big airport in Iowa to import wheat from China, on the basis that wheat costs less there. It just doesn't make any sense, given the abundance of natural gas in our country."*
>
> *"Madness," Stansberry's Investment Advisory, May 2006*

As everyone knows by now, we were exactly right about soaring gas production and the future of a business that was trying to import LNG to America.

We recommended shorting Cheniere Energy (NYSE: LNG, $140) back in May 2006 at around $40 per share. By 2008, it had fallen to about $2.00—a complete collapse. As we wrote at the time, "If you were trying to win a competition for the worst business idea, this one would be hard to beat."

But a funny thing happened on Cheniere Energy's Road to bankruptcy.

The company's founder and CEO had a complete change of heart. He realized, albeit a little late in the game, that the problem America faced wasn't a shortage of natural gas but a glut. The only way to solve this problem long term was to begin exporting massive quantities of natural gas via LNG.

In a case of real life being stranger than fiction, in 2009 the company completely made a mid-construction U-turn and, rather than building LNG import facilities, reverse engineered and rebuilt its facilities to handle LNG exports.

America loves a comeback story. Ever since Cheniere got on the right side of the most important trade in the world (the inevitable global domination by U.S. natural gas), the stock has basically moved in a straight line from $2.00 to $180, for a market cap today of $41 billion!

Cheniere is the largest LNG exporter from America, filling a crucial bottleneck in global energy markets. With revenues of $20 billion annually, the company is projected to earn about $13 per share in 2024

But Cheniere isn't going to dominate the global markets. It doesn't own any natural gas resources – it only owns the terminals.

Just imagine what's going to happen when entrepreneurs as savvy as the Rice brothers get involved in LNG. Remember, the people who built Cheniere knew so little about America's natural gas assets they were going to import gas to America.

Cheniere is the story of a monkey finding a dollar and thinking he's a banker.

The coming revolution is far bigger than Cheniere. America has more natural gas infrastructure than the rest of the world combined. America already produces more natural gas than any other country and has the capability to grow production faster too.

In December 2023, U.S. LNG exports set a new monthly record of 13.6 billion cubic feet, which is more than 20% of the world's current demand. This made America the single largest LNG exporter, surpassing Qatar. By 2028, with a series of new LNG export facilities coming online, America's export capacity will jump by another 80% to reach 25 billion cubic feet per day.

And so far, most of the major “frackers” have been excluded from the upside of America’s LNG export boom. That is, they have typically sold their gas at market rates here in the U.S. - allowing the LNG exporters to capture the spread between domestic and international gas prices.

In January 2024, EQT changed the game. It struck a gas supply deal with Texas LNG Brownsville LLC, a company building the Texas LNG project in the Port of Brownsville, located on the southern tip of Texas with access to shipping routes in the Gulf of Mexico. This wasn’t the typical LNG supply deal, where the owner of the most valuable asset – the gas itself – gets excluded from the upside. EQT became one of the first independent U.S. gas producers to sign a long-term supply agreement with an LNG exporter in the form of a “tolling agreement.” In this arrangement, EQT pays a “tolling fee” to the LNG terminal to liquefy its natural gas. After that, EQT is free to sell that LNG to the highest bidder on the international market - giving it the opportunity to capitalize on much higher gas prices overseas.

The Texas LNG project will begin coming online in phases starting in 2025. By 2028, it will reach a maximum capacity of 4 million tons per year (or about 200 billion cubic feet). EQT booked half of that capacity all to itself – 2 million tons per year, or roughly 100 billion cubic feet of gas production.

EQT’s CEO Toby Rice commented on the deal in EQT’s subsequent Q1 2024 earnings call in February, explaining:

> *“Our more integrated approach to LNG exposure compared with peers gives us direct connectivity to end users of our gas globally and we have seen strong interest from prospective international buyers.”*

EQT is slated to receive its first LNG from the Texas LNG plant in 2028. When that happens, everything will change.

We are of course talking about EQT.

EQT has grown production by 50% since the Rice Brothers took control in 2019. There’s no question the company can supply gas to the world for decades. EQT has 27 trillion cubic feet of proven natural gas reserves and controls 2 million acres of the Marcellus, the world’s richest natural gas field in America. In fact, the company’s assets have provided 10% of all U.S. natural gas production growth since 2005.

Meanwhile, global demand for LNG—especially American LNG—is soaring. Why?

Russia supplies Europe with 40% of its natural gas.

With this supply cut off, and much of the gas landlocked in Russia with no alternative shipping routes into other markets, global gas supplies have been squeezed. Prices for natural gas in Europe and Asia now consistently trade at $10 - $12 per MMbtu versus $2 - $3 in the U.S.

This has created a massive arbitrage opportunity for the producers and exporters of U.S. natural gas.

This enormous spread between international gas prices versus America has led to a huge shift in global supply. The U.S. LNG exporters have stepped in to fill the gap, with a record-setting year in 2023 that made America the world's largest LNG supplier.

An even larger opportunity is to replace coal internationally as the leading baseload power fuel.

As Europe is discovering, it isn't yet feasible to power an entire economy's electric grid with wind and solar power. The wind doesn't always blow, and the sun doesn't always shine. But simply replacing coal with natural gas (distributed as LNG) would dramatically reduce greenhouse gas emissions.

As Toby Rice explained in a 2022 conference call:

> *"Without incremental U.S. natural gas, the world is reverting to coal. In just the last 12 months, emissions associated with international coal consumption increased at a level that effectively wipes out all of the progress made by the United States in deploying wind and solar over the last 15 years. We will not be successful in addressing climate change without providing a scalable solution to international coal. That scalable solution is natural gas, and we are the ones that have it."*

EQT's plan is simple. Continue to increase production and forge strategic relationships with owners of LNG infrastructure to support global distribution. And most importantly, make deals that give EQT full control over the economic fate of its gas. Doing so will allow the company to capture far higher international prices for natural gas. This would vastly lower global emissions because it would take coal offline.

EQT has what the world needs most right now—virtually unlimited supplies of low-cost natural gas. In the short term (the next 3-5 years), these assets will be unlocked by new pipelines and new LNG terminals to supply Europe—especially Germany, Poland, and Lithuania—with reliable, long-term natural gas supplies of natural gas.

There's no more valuable strategic asset in America's effort to contain Putin's aggression than our natural gas supplies. And over the long-term (next 10-20 years), there's no other company better positioned to profit as the world takes coal offline. EQT's natural gas will be powering the grid – and electric cars – across America and around the world.

It's natural gas – not gasoline, that matters.

Over the longer term, EQT's efforts to become the world's "cleanest" energy company will show the path forward for our entire economy's energy needs.

Cars aren't going to run on gasoline for much longer. They can't run on solar power. Likewise, using solar and wind power exclusively for the power grid isn't feasible. As more cars depend on the grid for power, the amount of electricity consumption in the U.S. (and around the world) will soar.

What is the only clean, safe, and dependable way to supply that demand? American natural gas.

Which energy company will be America's (and possibly the world's) largest supplier of energy?

That will be EQT.

THE BIG SECRET ON WALL STREET

The Keystone

- A Matter of National Security
- How U.S. Defense Technology Will Supercharge the AI Industry

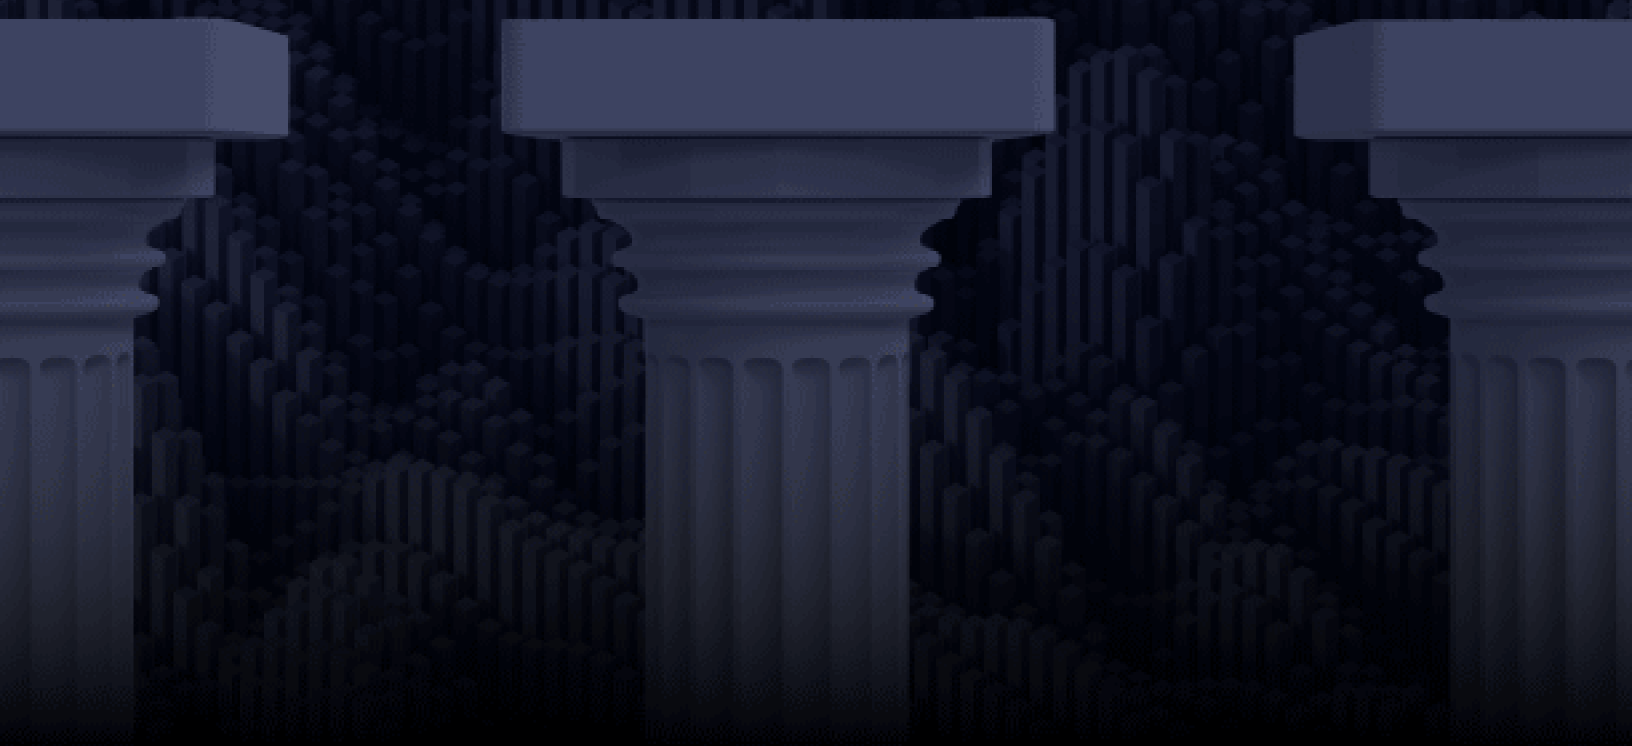

FROM THE DESK OF PORTER STANSBERRY

The Keystone

A Matter of National Security

How U.S. Defense Technology Will Supercharge the AI Industry

Quinnebago Outdoors is closed for good.

After two years of slow sales in his Panora, Iowa, sporting goods store – and around $25,000 in lost rental revenue for canoes and kayaks – store owner (and clever namesake) Quinten Pfeiffer auctioned off his inventory, and in November 2023 put his store up for sale.

There's nothing wrong with Quinten's business model. It's just that the Middle Raccoon River – where most of his customers swim, row, and float – has dried up.

It's pretty hard to paddle a canoe on mud.

Central Iowa is in the middle of a three-year-long drought – its worst in 20 years. Major tributaries, like the Raccoon and Skunk Rivers, are so low they're at wading level. The city of Des Moines is suffering "abnormally dry conditions," according to the U.S. Drought Monitor.

Unfortunately for Quinten and Quinnebago – and other Iowans' homes and businesses – the dry spell isn't likely to end anytime soon.

And it's not just because of the weather. The Drought Monitor reports that overall dryness in the state is improving… and in fact, October 2023 saw above-average rainfall.

So why are the rivers around west Des Moines still only ankle-deep?

For a surprising reason… ChatGPT.

Sometime in 2020, Microsoft quietly built a cluster of million-square-foot data centers in the cornfields of Iowa, in order to "train" the powerful artificial intelligence ("AI") program on human speech patterns and vocabulary. They were "pretty secretive on what they're doing out there," said Des Moines Mayor Steve Gaer. But eventually, residents realized that the company was siphoning off immense amounts of Iowa water to cool down ChatGPT's super-hot supercomputers.

Several rivers' worth, in fact.

It's no secret that the AI boom is energy-intensive. Since 2010, the processing power required to train AI went from doubling every 20 months, to doubling every six months. That's more than a 16-*trillion*-percent increase over the last 14 years alone… and it shows no signs of slowing down. Research institute Epoch AI warns that the

> *"amount of [computing power] developers use to train their systems is likely to continue increasing at its current accelerated rate."*

That means in one year, the amount of computing power required to train AI systems will increase by 300%… in three years, by 6,300%… and in five years, by a staggering 102,300%.

Along with that unimaginable amount of energy comes an unquenchable need for water to keep data processing centers from overheating.

And the more we rely on ChatGPT… and Bard… and DALL-E… and Midjourney… the more water we use. Just one plant – Microsoft's 102.5-acre ChatGPT processing center in Quincy, WA – diverts 121,000 gallons of river water *per day* for each of the three data servers on the campus.

It's not hard to guess where all the river water in Quinten Pfeiffer's neighborhood has gone.

As Virginia Tech researcher Abu Bakar Siddik told Futurism magazine, the "increased number of data centers that ramp up the water demand in the [Des Moines] region… could lead to high water stress in the region. Iowa can be disrupted by these events because there are already a high number of data centers compared to other states in the Midwest."

Microsoft used almost 1.7 billion gallons of additional water – increasing the company's global water consumption by about 34% – in 2022, largely due to the energy-intensive computer programs needed for its ChatGPT model. That's equivalent to over 2,500 Olympic-sized swimming pools. (Much of that water came from rivers in the already drought-burdened Des Moines area, where Microsoft had set up its data centers near cheaper power sources.) The same year, Google used *5.6 billion* gallons of extra water while "training" its Bard AI.

For perspective, every five to 50 questions you ask ChatGPT use up the equivalent of around a 16-ounce bottle of water. Right now, ChatGPT receives an average of 10 million queries per day (or 250,000 gallons of water). That's roughly 300 million queries (or around 7.5 million gallons of water) per month.

Research from the University of Washington shows that it costs around one gigawatt-hour of energy to answer these queries. One gigawatt-hour is roughly the same energy consumed by 33,000 households!

And that's for only one AI program. ChatGPT handles only about 60% of overall AI queries. And the field is still in its infancy.

The resources needed to sustain the AI boom – not just water, but energy itself – are staggering. And they are demands that, soon, we'll be unable to meet.

The Only Way to Generate Enough Power for AI

By 2027, *The New York Times* reported, AI servers are predicted to consume as much as 134 terawatt hours annually.

So, in less than three years, AI will have the same annual energy consumption as countries like Argentina, the Netherlands, and Sweden.

These energy demands are simply not sustainable.

Especially when AI adoption is at 1% of where we'll be in the next few years, according to industry insiders.

When AI adoption penetrates 10% or 20% of the market, we're looking at it consuming unprecedented amounts of energy.

Now, those on the frontier of the artificial-intelligence industry are aware of this problem, and they have started to sound the alarm.

> *"If you really want to make the biggest, most capable super intelligent system you can, you need high amounts of energy." - Sam Altman, CEO OpenAI*

> *"The world is actually headed for a really bad energy crisis because of AI unless we fix a few things." - Arijit Sengupta, founder Aible*

> *Elon Musk predicts that by 2045 the power demand in the U.S will have tripled from current levels – largely driven by AI's needs.*

They know that unless the insatiable energy demands of AI are met, the industry will never go mainstream.

And right now, there is no solution.

Fossil fuels are the primary energy source used to train and operate AI systems and could continue to meet the growing energy demands.

But that would require the woke Silicon Valley tech companies and progressive politicians to turn their backs on the religion of climate change.

And with the activists already up in arms about the environmental impact of AI, this is untenable to the ruling class.

So they'll argue that renewables like wind and solar should be used to meet the energy demands of AI. But the reality is that there is zero chance of these renewables producing enough low-cost energy to meet AI's needs.

Says a report from a New York City-based think-tank The Manhattan Institute: thinking that wind and solar can ever replace fossil fuels is nothing but an "exercise in magical thinking."

The projected energy consumption of artificial intelligence we discussed earlier is 134 terawatt hours (TWh) annually. To generate this amount of energy from wind power alone would require almost 17,000 wind turbines… taking up roughly 900 square miles, about 1.5 times the size of Houston – the fourth largest city in the U.S.

It's obvious that renewables like wind, solar, and hydro will never meet the energy demands of AI. And this has created a unique situation…

Bluntly, the economic, social, and geopolitical incentives behind AI are far too powerful to be stopped.

Even the Biden administration is not incompetent enough to halt the progress on A.I development and cede power to our rivals like China and Russia.

However, for AI development to continue and reach its full potential, vast amounts of energy will be required.

And there's only one viable solution: **nuclear energy**.

Just as the Keystone in an arch is required to hold the structure together, I believe the entire AI market will rely on nuclear power as a new Keystone technology.

Without this tech, I don't believe the industry will ever scale or achieve its full economic potential.

In short: nuclear energy is the future of the entire AI sector.

And – fascinatingly – right now, the nuclear opportunity we're recommending is being spearheaded, and funded, by the U.S. military.

No Snowflakes Allowed

When actual lives are on the line, it's interesting how priorities come into stark relief.

War, for instance, requires efficient energy that works. Full stop. Sorry, snowflakes. Grab a rifle or shut up.

During the height of the U.S. wars in Iraq and Afghanistan from 2001 to 2010, more than half of U.S. combat casualties were sustained during transport missions. And over 80% of these stemmed from demand for two critical battlefield resources: water and fuel.

Running a military requires a lot of energy. The Department of Defense (DoD) consumes 10 million gallons of fuel per day, and 30 TWh of electricity per year (one terawatt-hour is enough to power roughly 1 million households).

As military technology continues to advance, the energy demands of the modern battlefield will only increase. That spells more opportunities for the enemy to strike vulnerable fuel supply chains, leading to more American lives lost. To address this growing vulnerability, the DoD established the Task Force on Energy Systems for

Forward/Remote Operating Bases to find a solution. In August 2016, the task force released a report detailing its findings.

The report began by crossing out the solutions that don't work:

> *"The study found alternative energy sources, such as wind, tidal, solar, and other sources, were unlikely to comprehensively meet current or future energy demands for forward operating bases, remote operating bases, and expeditionary forces."*

Let's pause for a moment to appreciate the irony here. While U.S. politicians squander trillions of taxpayer dollars trying to overhaul America's formerly robust electric grid with unreliable wind and solar power, the DoD is running in the opposite direction.

To address the needs of military commanders tasked with winning on the battlefield with a minimal loss of life, snowflake economics and feel-good fantasies like solar and wind power need not apply.

That's how the task force settled on the most reliable, high-density energy known to man: nuclear energy.

Nuclear power is one of mankind's most remarkable achievements – offering a virtually limitless source of reliable, cheap, carbon-free baseload power. If environmentalists were actually moral scientists, using technology to build a better life for more people, they would be pounding the table on nuclear power. That they abhor nuclear power above all other solutions tells you all you should need to know about their real purposes.

They aren't saints. They are Nazis, determined to end human civilization as we know it.

Sure, it's a hobby horse for us – calling environmentalists Nazis. But what wouldcause the deaths of more people? The Nazis, who are largely responsible for World War II, and who are definitely responsible for murdering millions of Jews, caused the deaths of something like 30 million people. If the environmentalists could end coal-fired electricity tomorrow, upon which most humans on this planet depend?

Billions would die.

Nuclear power starts with the uranium-235 isotope. Scientists learned to "split" this atom in the 1940s through nuclear fission. The fission reaction unleashes unimaginably larger (1.5-2.5 million times more) amounts of energy per unit of mass compared to coal, oil, or natural gas. The fission of a 10-gram (a peanut weighs about a gram) uranium pellet releases as much energy as burning 4,350 gallons of oil... 22 tons of coal... or 590,000 cubic feet of natural gas!

For the DoD, a pebble or so of uranium-235 could replace thousands of fuel-hauling vehicle convoys, potentially saving the lives of countless American troops. It could also be used to power water purification and recycling, and other energy-intensive battlefield requirements. And of course, the civilian applications of this kind of technology – though less imminent – are beyond mind boggling.

So, why aren't we using nuclear energy everywhere?

Here's the problem – this powerful, efficient energy source is usually chained to unwieldy, giant nuclear power plants that take 10 to 15 years to build and billions of dollars of investment.

Building a full-sized 500 megawatt (MW) nuclear power plant in a battle zone isn't an option. But what could make sense is harnessing nuclear energy at 1/100th of that scale, providing power to supply the roughly five MW required to run the forward operating bases (FOB).

FOBs are small, makeshift military bases used in areas where a physical presence is needed, but where a full-scale military base is impractical. For example, during the war in Afghanistan, the U.S. military built FOBs in areas staffed by a few dozen troops on an isolated mountainside. A reactor to support an FOB would need to be able to be deployed by rail, truck, or cargo plane, and small enough to fit inside a 20-by-20-foot shipping container.

The 2016 task force concluded that such a reactor design was possible, and DoD's Strategic Capabilities Office moved to the next phase – building a prototype.

Project Pele

Project Pele brought together an alphabet soup of government agencies, including the Department of Energy (DoE), NASA, the Army Corps of Engineers (USACE), the Nuclear Regulatory Commission (NRC), and the National Nuclear Security Administration (NNSA). The project's aim: develop a nuclear microreactor for deployment "by road, rail, aircraft, or sea" that was also capable of "quickly being brought on land" and was "inherently safe." Success would be "a strategic game-changer for the United States, both for the DoD and for the commercial sector," according to Project Pele manager Jeff Waksman.

To make this ambitious plan a reality, the DoD enlisted help from the private sector. In March 2020, the DoD launched a two-year design competition for a prototype of the Project Pele microreactor, soliciting bids from a group of top nuclear engineering design firms.

In June 2022, the DoD selected a prototype developed by a public company that's poised to reap a windfall by developing the next phase of nuclear power – one that could revolutionize global electricity production around the world.

Before we talk about the massive new opportunity, let's review this company's highly profitable core business that offers the ultimate recession-proof play. After all, there's one trend we can count on through thick and thin… the growth in the U.S. military budget.

20,000 Leagues Under The Sea

The company tasked with building America's first commercial microreactor also built the world's first portable nuclear reactor… 70 years ago.

USS NAUTILUS Launching Ceremony
January 21, 1954

Just nine years after testing the first atomic bomb, America harnessed the awesome power of nuclear energy in a portable underwater reactor on board the U.S.S. Nautilus – the world's first nuclear submarine, launched on January 21, 1954 (and named after Captain Nemo's famous science-fiction submarine in *20,000 Leagues Under the Sea*).

Before the revolution in naval war unleashed by the Nautilus, submarines ran on diesel-electric power trains. These vessels were powered by a large bank of lead-acid batteries, which were charged by diesel engines. The batteries lasted only about two to three days before needing a recharge. The diesel engines that charged the batteries required oxygen – which meant that when the batteries depleted, submarines would need to surface in order to draw atmospheric oxygen. Going up to the surface, and emitting hot exhaust fumes, is the last thing that submarines should do if they want to remain undetected.

The Nautilus's S2W Thermal Nuclear Reactor reactor solved these problems. It required no oxygen, generated no external waste, and could travel an almost unthinkable 62,000 miles (that's like circling the earth twice, with 10,000 miles to spare) before needing to replenish its nuclear fuel. Food supplies, rather than fuel, became the limiting factor for submarine voyages. What's more, the nuclear engine generated a monstrous 13,400 horsepower, making it significantly faster and more maneuverable than its diesel-electric peers.

On May 10, 1954, the Nautilus made naval history when it traveled 1,400 miles from Connecticut to Puerto Rico, fully submerged, in less than 90 hours – an evening short of four days. This shattered records for the longest and fastest submarine cruise, and it rendered obsolete the entire playbook on anti-submarine warfare tactics developed throughout World War II. It also set the stage for America's undisputed naval supremacy for the next 70 years.

The importance of controlling the seas – the conduit for 80% of global trade – has been known since ancient Greek statesman Themistocles famously declared, "He who controls the sea controls everything."

The company that designed and built the components for Nautilus's nuclear reactor was **BWX Technologies (NYSE: BWXT)**. Since then, BWXT has cemented itself as the key supplier of reactor design, components and fuel for America's nuclear navy.

Mr. Monopoly

U.S. defense spending grows even – or especially – during recessions, which makes this defense supplier stalwart a great investment for our current **Minsky moment**.

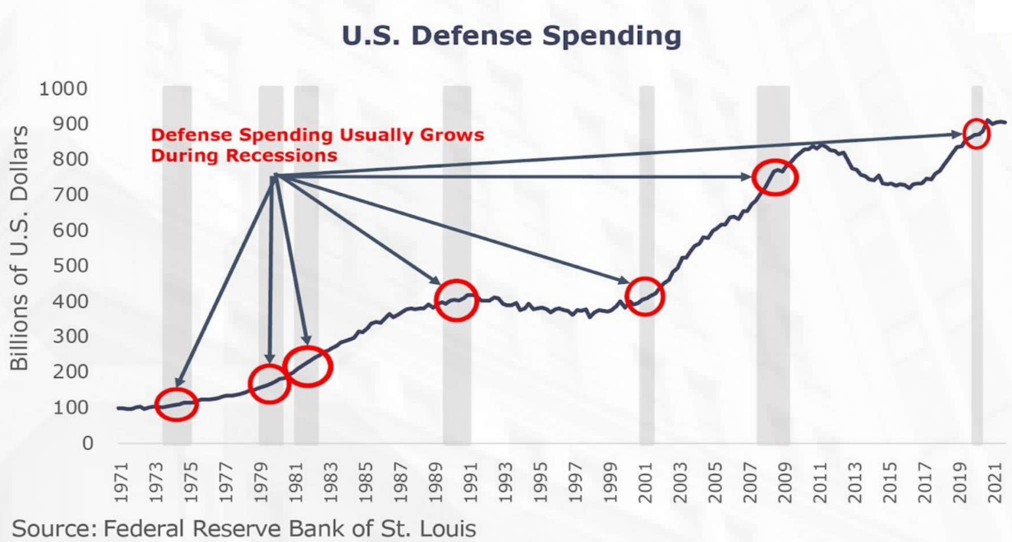

BWXT has been at the forefront of nuclear technology since the birth of the industry. Today it operates four main nuclear business units in the U.S. (BWXT Power, BWXT Nuclear Energy, BWXT Nuclear Operations Group, and BWXT Technical Services Group), as well as the only two commercial plants in the U.S. that process uranium.

Most importantly, BWXT is the sole manufacturer of nuclear reactors and fuel for U.S. military aircrafts and submarines. It's also one of only two providers licensed to store and process HEU (highly-enriched uranium) for these reactors.

The government sector accounts for 80% of BWXT's revenue. And nearly all of its deals with Uncle Sam are carried out via long-term contracts, resulting in a very stable and predictable business and revenue flow.

As a monopoly manufacturer, BWXT can set prices (within reason), and lock in steady profit margins. If a project incurs unexpected cost overruns, BWXT can charge the Navy back fees to make sure it hits its target profit margins. That's how BWXT posts stable profit margins, with very little exposure to swings in the economy from recessions, inflation, or other external factors.

Over the last five years, the economy suffered through a devastating pandemic and economic shutdown, followed by the hottest inflation in the last 40 years. During one of the most turbulent macroeconomic periods in U.S. history, BWXT's business has chugged along with remarkable stability with profit margins:

Source: Bloomberg

The business also provides a clear line of sight into the future, based on the full slate of projects BWXT has lined up...

BWXT Is The Navy's "Main Squeeze"

With around 300 active, deployed ships, the U.S. Navy is not the largest in the world by ship count. It is, however, the most powerful naval force thanks to its unmatched fleet of nuclear-powered vessels, including 53 attack submarines, 18 strategic submarines equipped with nuclear warheads, and 11 aircraft carriers.

BWXT is responsible for powering the Navy's submarines and its aircraft carriers. It delivers nearly all of the mechanical equipment in the engine room for the Navy for their nuclear platforms. Below is the fleet of Naval nuclear platforms, as well as the status (number of orders left to fulfill) and the contribution value of each ship when serviced or refueled.

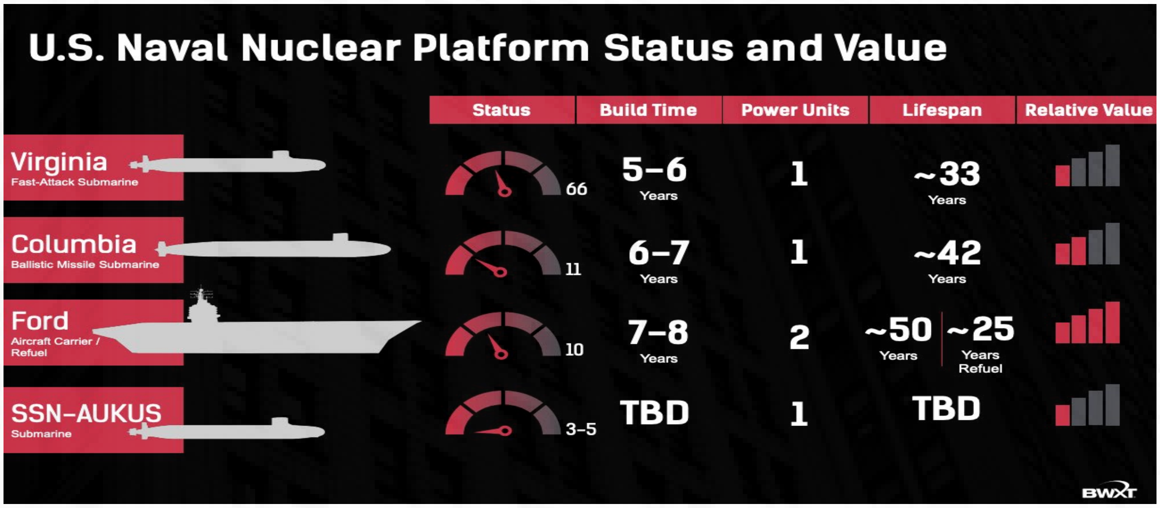

Source: BWXT 2024 Investor Day Presentation

The Virginia, Columbia, and SSN-AUKUS submarines, as well as the Ford aircraft carrier, depend on BWXT, allowing for a long runway of growth for the company over the next five to eight years.

The Virginia and Columbia have life-of-ship power units (meaning they only need one reactor for the life of the vessel), while the Ford has half-life-of-ship power that need to be replaced once after 25 years. The Ford, as seen above in the chart, has the highest relative value followed by the Columbia.

In recent years, the U.S. and its allies have committed to bulking up their nuclear naval fleets in response to growing geopolitical tensions with China and Russia. This includes the trilateral agreement between the U.S., UK, and Australia – known as AUKUS. Starting in 2023, the U.S. and UK began training Australian sailors on the technology and operations of nuclear submarines. The training has set the stage for deliveries of U.S. Virginia-class submarines to Australia in the early 2030s. In the 2040s, the three countries will collaborate to build the next generation of nuclear submarines.

As part of this agreement, BWXT will supply reactor components for 3 – 5 Virginia class submarines to be delivered to Australia starting in the 2030s. BWXT will also play a key role in designing and supplying reactor components for the next-generation submarines developed by AUKUS in the 2040s. BWXT is currently negotiating the pricing for this work.

These contracts provide BWXT visibility into future orders, which BWXT receives two years prior to the shipbuilders' receiving their order. Below are the scheduled programs which BWXT released at its 2024 Investor Day.

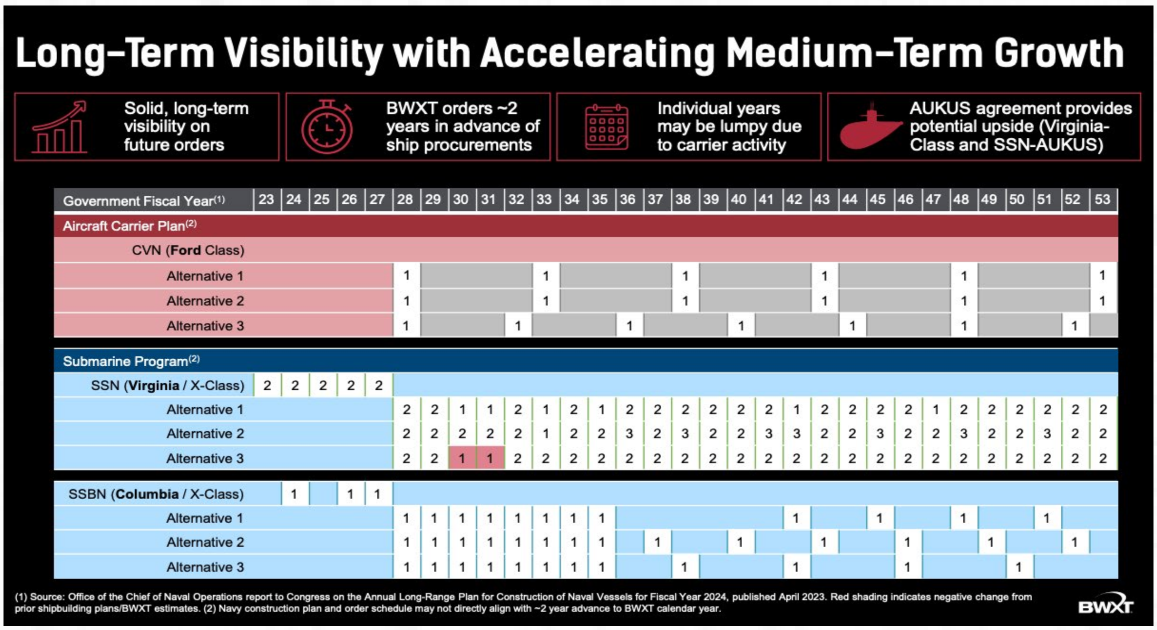

Long-Term Visibility with Accelerating Medium-Term Growth

- Solid, long-term visibility on future orders
- BWXT orders ~2 years in advance of ship procurements
- Individual years may be lumpy due to carrier activity
- AUKUS agreement provides potential upside (Virginia-Class and SSN-AUKUS)

Government Fiscal Year[1]	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53
Aircraft Carrier Plan[2]																															
CVN (**Ford** Class)																															
Alternative 1						1					1					1					1					1					1
Alternative 2						1					1					1					1					1					1
Alternative 3						1				1				1				1				1				1				1	
Submarine Program[2]																															
SSN (**Virginia** / X-Class)	2	2	2	2	2																										
Alternative 1						2	2	1	1	2	1	2	1	2	2	2	2	2	2	1	2	2	2	2	1	2	2	2	2	2	2
Alternative 2						2	2	2	2	2	1	2	2	3	2	3	2	2	3	3	2	2	3	2	2	3	2	2	3	2	2
Alternative 3						2	2	1	1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
SSBN (**Columbia** / X-Class)		1		1	1																										
Alternative 1						1	1	1	1	1	1	1	1							1			1			1			1		
Alternative 2						1	1	1	1	1	1	1	1		1			1			1			1			1			1	
Alternative 3						1	1	1	1	1	1	1	1			1				1				1				1			

(1) Source: Office of the Chief of Naval Operations report to Congress on the Annual Long-Range Plan for Construction of Naval Vessels for Fiscal Year 2024, published April 2023. Red shading indicates negative change from prior shipbuilding plans/BWXT estimates. (2) Navy construction plan and order schedule may not directly align with ~2 year advance to BWXT calendar year.

BWXT

Source: BWXT 2024 Investor Day Presentation

Orders from the Ford and the Columbia ramp up in 2028, which directly correlates with future revenue and earnings for BWXT. With the Ford having the greatest impact on BWXT's top and bottom lines, and the Colombia receiving one expected order per year between 2026 and 2035, BWXT's roadmap over the next decade is promising, and poised to translate to growth in earnings.

BWXT's margins increase over the time of a contract as it typically realizes cost underruns and optimizes the manufacturing costs over the lifespan of naval nuclear-reactor contracts.

Very Stable Genius

The predictability and stability of BWXT's business is unmatched. It has a near-lock on supplying the Navy with critical nuclear inputs. And demand for new carriers and submarines is based largely on the retirement of the existing fleet, so it's easy to map out what the business will do next. And it's growing apace...

Revenue grew from $1.6 billion in 2016 to $2.5 billion over the last 12 months, while earnings rose from $184 million to $263 million. The company generates roughly $260 million in free cash flow ("FCF") and spends between $100 and $200 million in capital expenditures ("capex") each year. This FCF goes toward a steady buyback program, which has reduced the outstanding share count by 14%, from 105 million in 2016 to 91 million by the end of 2023. (Share buybacks are a tax-efficient way of returning value to shareholders.)

The combination of growing net income and a falling share count has boosted earnings per share ("EPS") from $1.79 in 2016 to $2.78 in 2023.

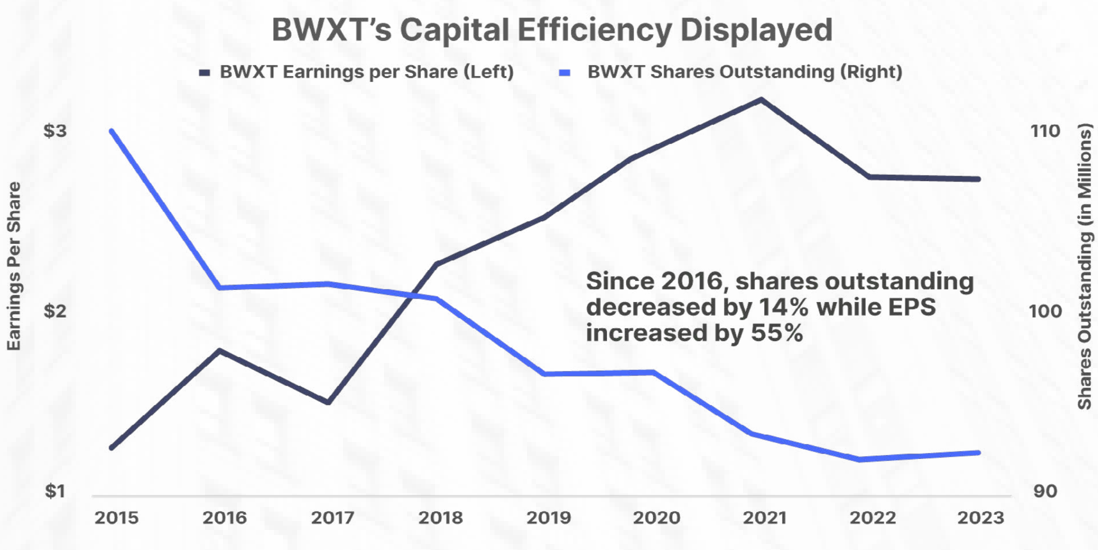

Source: Bloomberg

The company's balance sheet is conservatively managed, with $1.2 billion in long-term debt, supported by roughly $400 million in annual operating income. So BWXT offers investors a very stable, recession-proof business that's well-positioned to thrive in an uncertain economy.

But by far, the most exciting part of this story will come from BWXT's work on the Project Pele reactor, and its huge upside potential. In June 2023, BWXT won the design competition for the Project Pele reactor prototype, and received a $300 million contract to build a full-scale version.

BWXT's vertically integrated approach gives it a leg up over companies that depend on other countries – like TerraPower, Bill Gates' pet nuclear project, which also features more compact nuclear reactors but unfortunately relies on Russian-produced HALEU fuel. In 2021, TerraPower announced that it will halt operations for two years due to the Ukraine invasion. Now TerraPower isn't expected to generate power until 2030 while fuel supply remains a concern.

BWXT is pioneering "bring-your-own-nuclear" energy, which could rip up the playbook of energy as we know it today.

The Project Pele microreactor received an initial contract of $300 million to deliver up to 5 megawatts of electrical power. But the DoD uses 30 terawatts – that is, 30 million megawatts – of electricity per year, opening the door to enormous revenue growth.

The DoD Strategic Capabilities Office (SCO) partnered with BWXT to build the first advanced nuclear reactor of its kind and further engrains BWXT's role as a critical military supplier, while bolstering BWXT's position as a nuclear power pioneer.

Assuming the first full-sized prototype meets all required specifications (to be determined later in 2024), the initial use case for the DoD involves deployment of these reactors in forward operating bases. There's additional potential for these reactors to be deployed in disaster-relief zones, both domestically and abroad. The reactor could also serve as a "pathfinder" for commercial adoption of such technologies, DoD said.

Explained SCO director Jay Dryer:

> *"The DoD has a long history of driving American innovation, with nuclear power being one of many prominent examples. Project Pele is an exciting opportunity to advance energy resilience and reduce carbon emissions while also helping to shape safety and non-proliferation standards for advanced reactors around the world."*

And maybe beyond earth too. NASA is in the process of developing a human base on the moon, as a precursor to manned Mars exploration. NASA chose BWXT to develop the nuclear-based propulsion systems for that program, called Artemis.

Since the beginning of Europe's energy crisis, it's been obvious that the world will eventually vastly increase its use of nuclear power. Technology and humanityconsistently evolve toward more dense forms of energy. And with each evolution of power technology, human wealth grows exponentially.

The next 50 years will almost certainly be the age of nuclear power. There are virtually unlimited applications for the small, safe, and portable reactors that BWXT builds – and we believe the enormous energy demands of AI will be foregrounded.

For now though, BWXT offers a stable business model that's recession-resistant, a durable competitive advantage, and the upside kicker of advancing small modular nuclear reactors, first to the military, and potentially to the world.

Just how big of an upside kicker?

The first reactor will cost $300 million. But BWXT will benefit from economies of scale that cut costs over time. A rough estimate suggests BWXT could get costs down to anywhere between $100 and $200 million. If BWXT sells 150 reactors at $100 million to $200 million apiece, that translates into $15 billion to $25 billion in new revenue. For a company with a current market capitalization of $9 billion, that's a powerful upside catalyst.

BWXT plans to complete construction and deliver the first full-sized Pele reactor sometime in 2025. It will then undergo a series of tests at the Idaho National Laboratory to ensure it meets DoD specifications. Until these test results come in, we can't know exactly how much future demand will exist, if any. But if the reactor does meet the ambitious DoD requirements, then the sky's the limit.

In the meantime, BWXT's core business offers the ultimate safe haven against the economic storm clouds gathering on the horizon. With years of backlogged demand, and a dominant competitive position generating rock-solid profit margins, we feel comfortable recommending this stock based on its existing business today. Plus, we're getting plenty of potential upside from Project Pele.

PART 4: HOLDING IT TOGETHER

Energy isn't enough. As we've just discussed, a surge of new power sources – coal, gas, and nuclear – are coming online to power the Parallel Processing Revolution. That means we need data centers... telecommunications... factories... fiber optic networks... power transmission lines... satellites... manufacturing... transportation... and the list goes on. In other words, we desperately need a multitrillion-dollar overhaul of America's aging energy grid.

This overhaul is already underway, thanks to a record injection of government stimulus money, but it will take years and cost billions of dollars to finish. And one company could play a critical role in the rebuild of the grid: **Atkore (NYSE: ATKR)**, a leading manufacturer of electrical components used in construction and manufacturing.

In our final report, **The Big Long**, we'll examine Atkore's competitive advantages and show how it holds the Parallel Processing Revolution together in the midst of exponential growth.

Porter & Co.

THE BIG SECRET ON WALL STREET

The Big Long

- How to Profit From America's Decades-Long Infrastructure Boom
- This Company Unlocks Premium Pricing Power From Selling Commodities

FROM THE DESK OF PORTER STANSBERRY

The Big Long

How to Profit From America's Decades-Long Infrastructure Boom

This Company Unlocks Premium Pricing Power From Selling Commodities

"I can't f***ing believe this is allowed."

... so muttered fund manager Steve Eisman "a thousand times" as he uncovered layer after layer of fraud and deception permeating America's housing market between 2006 and 2007.

Eisman was no stranger to financial chicanery. He got his start on Wall Street analyzing the shady underbelly of American finance – subprime-mortgage lenders, a business model built on making home loans to borrowers at high risk of not repaying their debts.

Eisman made a name for himself by doing something most Wall Street analysts shy away from – advising his clients to "sell" instead of "buy." And he did it with no reservations. He once walked into the middle of the trading floor at Oppenheimer & Company (where he worked as an analyst) and made an announcement: "The following eight stocks are going to zero." He listed the eight companies, all of which eventually went bust.

The problem was, Eisman's cynicism didn't help drum up much business for the investment banking firm he worked for. They were in the business of helping corporations raise capital – not bash their stock prices. So in 2004, Eisman launched his own hedge fund, FrontPoint Partners. He was then free to bet against bad business models instead of simply writing reports about their inevitable demise. His previous experience with subprime lenders in the 1990s made him the perfect candidate to bet against what became the mother of all manias: the subprime-housing bubble of the early 2000s.

When Eisman was analyzing subprime lenders in the 1990s, they were a tiny piece of the overall mortgage market. Back then, $30 billion a year in new subprime loans was a big deal. But by 2000, that number surged to $130 billion, before ballooning to $625 billion in 2005. Eisman knew that a lot of bad loans were being made, and that the companies making them would eventually go belly-up.

Eisman bet against the subprime-mortgage craze by shorting stocks like New Century Financial – one of the leading originators of subprime mortgages. But the reward did not necessarily outweigh the risk. The most Eisman could make was 100% if a shorted stock went to zero. Even those profits got eaten away, as short sellers had to pay a 20% dividend to the shareholders they borrowed stock from in order to bet against it. Plus, short sellers had to pay a 12% interest rate along the way. This meant Eisman had to shell out $32 million per year for the privilege of holding a $100 million short position. Meanwhile, if the shares increased in value, he faced the prospect of unlimited losses.

Then in February 2006, a Deutsche Bank trader named Greg Lippmann introduced Eisman to a different way of betting against the subprime-mortgage market. Credit default swaps ("CDS") provided a form of insurance against a subprime-mortgage meltdown. They offered the ability to bet directly against the subprime loans being issued, instead of against the companies that were making the loans.

The upshot was that buying CDS contracts could produce returns of as much as 1,000% on an original investment. The lopsided return proposition stemmed from the fact that the banks selling the insurance contracts had grossly mispriced the odds of a subprime bust. Eisman was elated, describing the conversation he had with trader Lippmann:

> *"When he walked in and said you can make money shorting subprime paper [from buying CDS contracts], it was like putting a naked supermodel in front of me."*

Over the next 18 months, Eisman and his crew explored the underbelly of Wall Street's subprime machinery. He learned how Wall Street, with help from the ratings agencies, packaged low-grade mortgages into AAA-rated fixed-income securities known as mortgage-backed securities ("MBS"). And how even lower-quality bonds were packaged into collateralized debt obligations ("CDO"). The credit default swap contracts he purchased provided a way to profit if these loans went bust.

The deeper Eisman dug, the more incredulous he became. For instance, he learned that bonds were being made from mortgages, known as "no-doc" loans, for which no documentation was required about the borrower's job history, income, or assets. As Eisman explains...

> *"The first time I realized how bad it was, was when I said to Lippmann, 'Send me a list of the 2006 deals with high no-doc loans... I figured Lippmann was going to send me deals that had 20% no docs. He sent us a list and none of them had less than 50%."*

That meant that half of all mortgages were being granted to borrowers who provided no proof of their ability to repay the loans. This lack of oversight is how Wall Street's subprime machinery allowed borrowers to take out more money than they could ever repay. Eisman discovered loans like one made to a strawberry picker in Bakersfield, California, who had income of $14,000 per year, and yet was able to secure a loan to purchase a $724,000 house.

Early on in the housing boom, these loans didn't go bust because borrowers paid rock-bottom interest rates for a short period – before the rates contractually reset at higher levels a few years down the road. Before then, though, with home prices rising, borrowers could refinance for more than the original value of the property, receiving a cash infusion that enabled them to keep making their monthly payments. It was a perfect Ponzi scheme – homeowners, mortgage issuers, and investment bankers would all get rich if home prices kept moving higher.

But Eisman figured it was only a matter of time before home prices peaked, and interest rates reset at higher levels, triggering a wave of defaults. It was obvious, he thought – why could no one else see what's coming?

What kept the bubble from bursting was the complicity of the ratings agencies.

Eisman discovered how corrupt the ratings agencies had become in facilitating the sale of these subprime mortgages through mortgage-backed securities. The agencies rated a substantial portion of bonds as AAA, or as good as U.S. Treasuries, based on a series of what proved to be ludicrous assumptions they plugged into their ratings models.

After personally meeting with analysts at the ratings firms, Eisman learned that the models used to rate bonds assumed borrowers' ability to pay back the loans would be unaffected by their mortgage rates resetting at higher levels. He also learned that the models had no input for a drop in home prices. They assumed, like many on Wall Street, that U.S. housing prices couldn't go down. Eisman was shocked by each new revelation, as he recalled:

> *"I cannot f***ing believe this is allowed. I must have said that one thousand times."*

Along the way, in typical Eisman fashion, he didn't hesitate to express his views on the situation. During a financial conference in Hong Kong, the chairman of HSBC bank claimed that the losses in his bank's subprime portfolio would be "contained." Eisman stood up from the audience and responded, "You don't actually believe that, do you? Because your whole book is f****d."

Eisman couldn't believe how even the supposedly "smart money" sitting at the top of the global financial system was sleepwalking into disaster. And he grew giddy about the prospect of cashing in on their ignorance. As he explained the situation to a coworker, he said:

> *"It's a gold mine. And nobody else knows about it."*

By January 2007, Eisman (pictured below) had placed an all-in bet on what he expected to be a housing armageddon. His fund had purchased $550 million in credit default swaps against subprime-mortgage bonds, believing that the bonds would quickly lose value. By June 2007, subprime-mortgage bonds began selling off slowly, and then suddenly. FrontPoint's positions began moving up in value by hundreds of thousands of dollars each day, and then by millions.

Over the following year, these positions produced a windfall for his fund investors as the subprime-bond market entered into freefall. Eisman's fund ballooned to $1.5 billion, netting his investors $1 billion in profits.

Today, Eisman is more optimistic about the future of the U.S. economy – going long more often than he shorts. In particular, there's one major theme that he's betting heavily on, which we'll dive into in this issue.

Eisman's Next Act: The Big Long

In a June 2023 interview on the Bloomberg *Odd Lots* podcast, Eisman made the case for what we're calling the Big Long, in reference to the 2015 *Big Short* film that profiled a group of fund managers, with Eisman as a major character, who profited from the collapse of the housing market in 2007-2008.

Eisman believes there are huge profits to be made from the record influx of capital reshaping and revitalizing America's industrial infrastructure over the next decade. This includes $2 trillion in stimulus spending from the federal government, plus trillions more from the private sector.

The first big theme driving this trend is the rise of onshoring, which involves U.S. manufacturers bringing their overseas operations back home. This trend began gathering momentum in 2018 after the Trump administration imposed a series of tariffs and import duties that upped the cost of outsourcing U.S. manufacturing to China.

The COVID-19 outbreak and economic turmoil that resulted massively accelerated this trend. The pandemic-driven supply-chain disruptions revealed the extreme reliance U.S. corporations had on overseas suppliers, for everything from computer chips to pharmaceuticals.

As a result, U.S. companies across many industries are pouring record amounts of capital into domestic manufacturing facilities. In the last three years alone, $500 billion of investment has gone into new U.S. factory construction, compared with about a total of $200 billion over the previous three-year period:

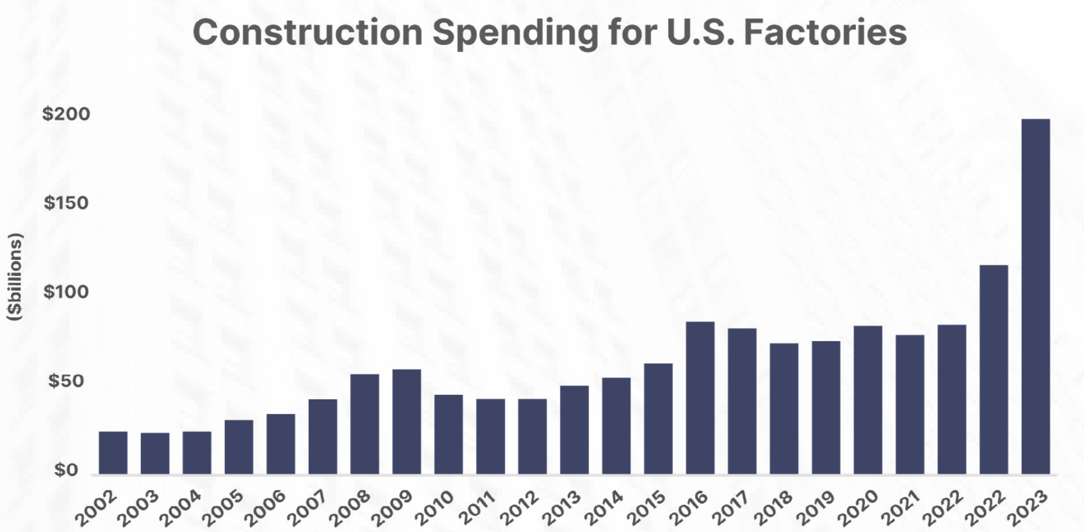

Source: Census Bureau November and December 2023 are estimates

Within this manufacturing bonanza, one area in particular is benefitting the most – investment into computers, electronics, and electrical power transmission:

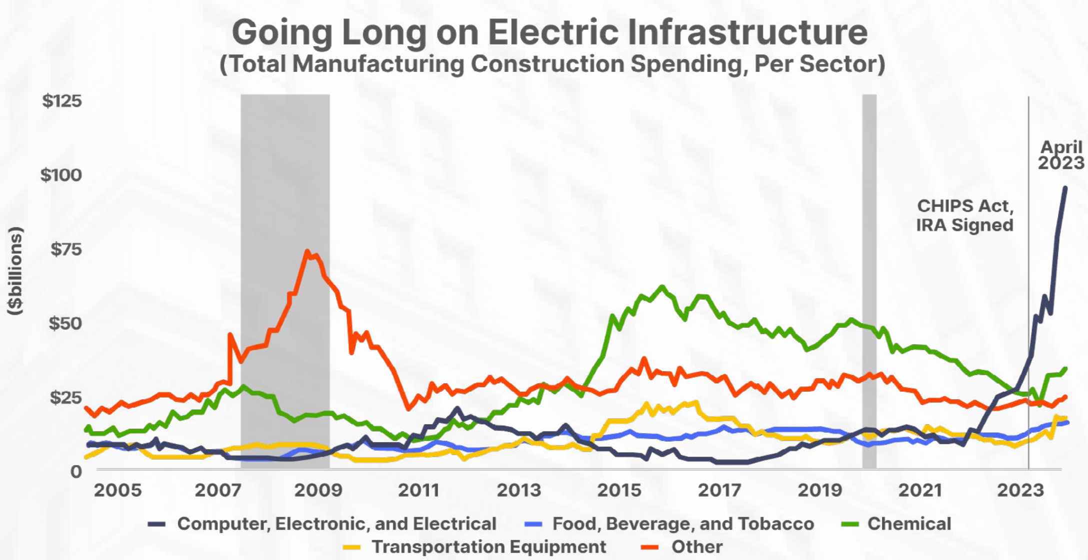

Three major factors are driving the building boom in this sector:

1. The rise of electrified transport (i.e., electric vehicles and charging stations) and "green" power generation (i.e., solar and wind power).
2. The artificial-intelligence ("AI") revolution that's fueling a boom in U.S. semiconductor manufacturing and data-storage centers.
3. The dire need to overhaul America's flailing electric grid in response to the massive new electrical infrastructure needs from the two trends noted above.

We'll dive deeper into each of these three segments later. For now, the key is that all of these trends are being turbocharged by $2 trillion in federal grants and incentive programs launched in 2021 and 2022. These include the $280 billion CHIPS and Science Act ("CHIPS"), the $579 billion Inflation Reduction Act ("IRA"), and the $1.2 billion Infrastructure Investment and Jobs Act ("IIJA"):

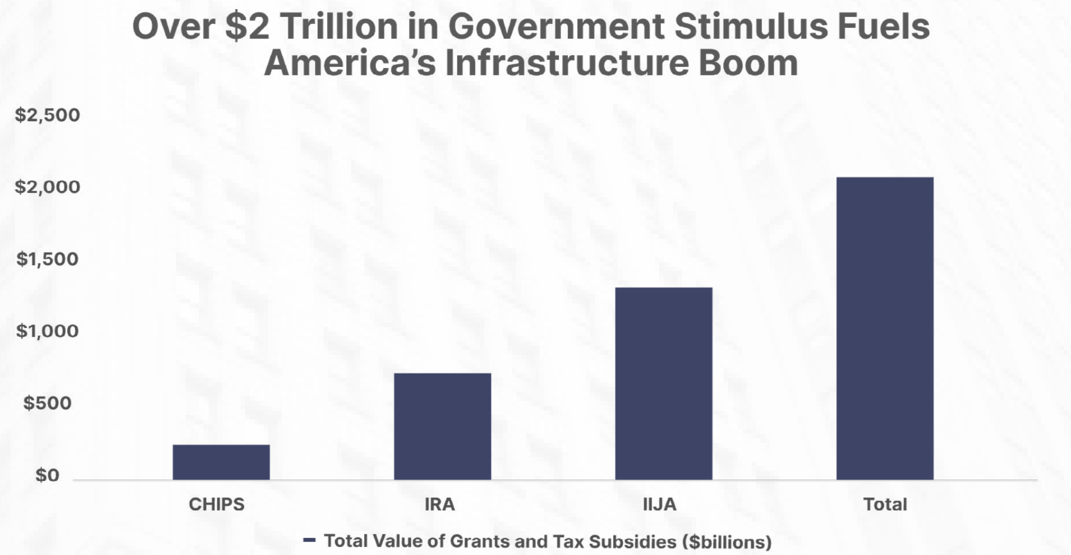

Source: Bloomberg

Together, these three pieces of legislation represent the largest infrastructure stimulus program in American history. And we're still in the very early stages of this infrastructure bonanza, with most of this money still slated to be allocated over the next decade.

In this issue, we're recommending our best idea for capitalizing on this decades-long infrastructure boom.

Selling the Components That Make It All Possible

The company we're recommending in this issue is **Atkore (NYSE: ATKR)**, a leading manufacturer of electrical components used in construction and manufacturing.

The company separates its business into two segments: Electrical, and Safety & Infrastructure. In 2023, Atkore generated $3.5 billion in revenue: $2.7 billion from its Electrical segment and $844 million in Safety & Infrastructure.

In Electrical, Atkore manufactures things like power systems, breaker boxes, electrical wiring, and plastic conduits, which is the protective tubing used to insulate electrical wiring. It also sells electrical mounting systems and installation accessories used for securing electrical products to a building's structure. Finally, this segment also includes fiber-optic cable and conduit and mounting systems for high-speed internet and telecommunications.

In Safety & Infrastructure, Atkore manufactures metal framing, fittings, and mechanical tubing used for support structures and conduits. It also sells perimeter security products, like chain-link fences and barbed wire, as well as cable management systems.

The graphic below shows a snapshot of the many different products it supplies across all of these different applications:

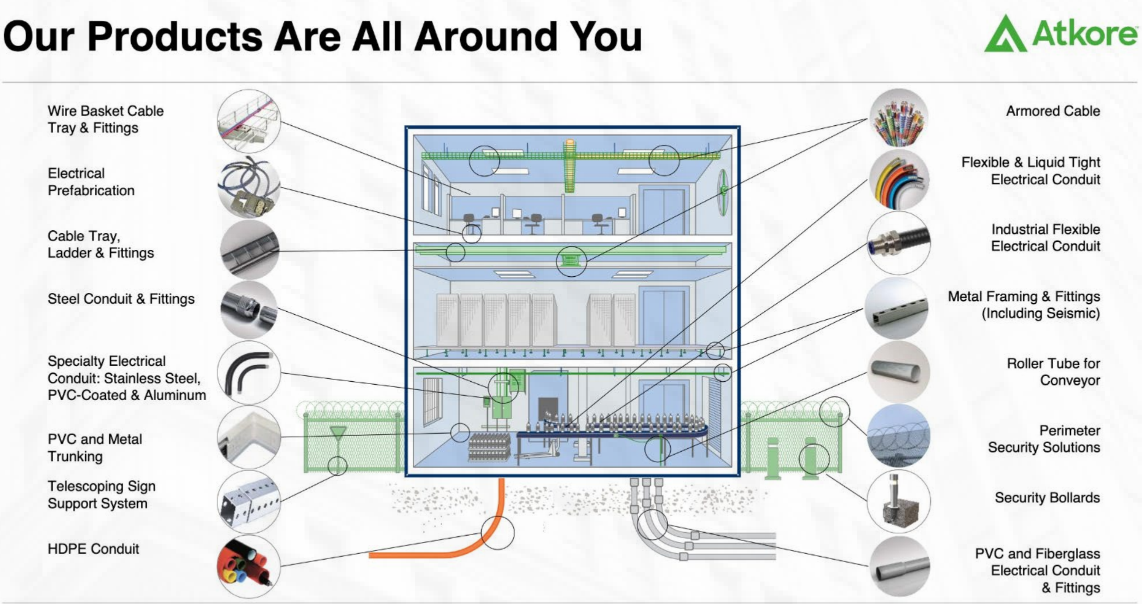

It's important to note that many of Atkore's Safety & Infrastructure products are used for electrical infrastructure. This includes things like wire-basket cable trays and steel tube conduit, both of which are used to support and protect electrical cables. This overlap in terms of end market consumption means that over 90% of Atkore's sales are used for electrical infrastructure:

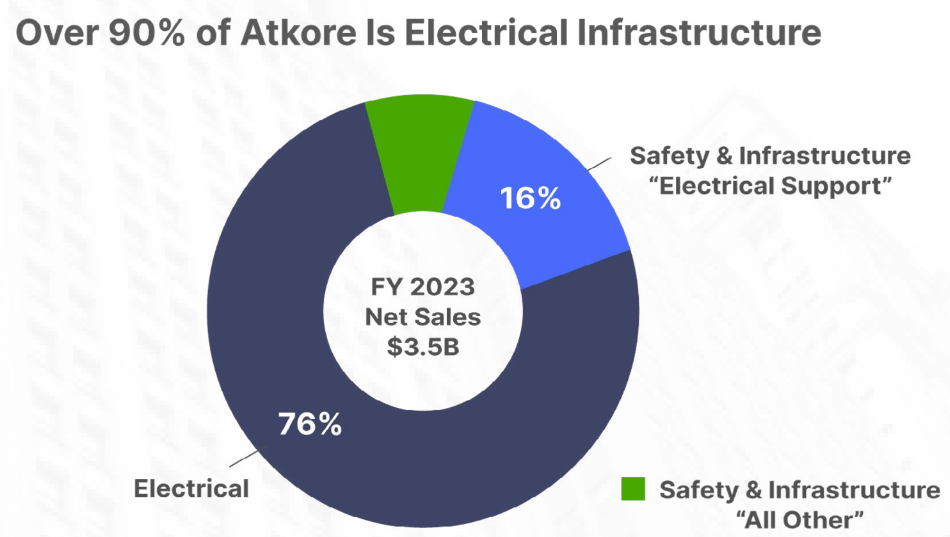

Given Atkore's exposure to the fastest-growing segment of U.S. infrastructure spending, business has boomed in recent years. Revenue doubled from $1.9 billion in the 2019 fiscal year (which ends September 30) to $3.9 billion by 2022, before falling to $3.5 billion in 2023.

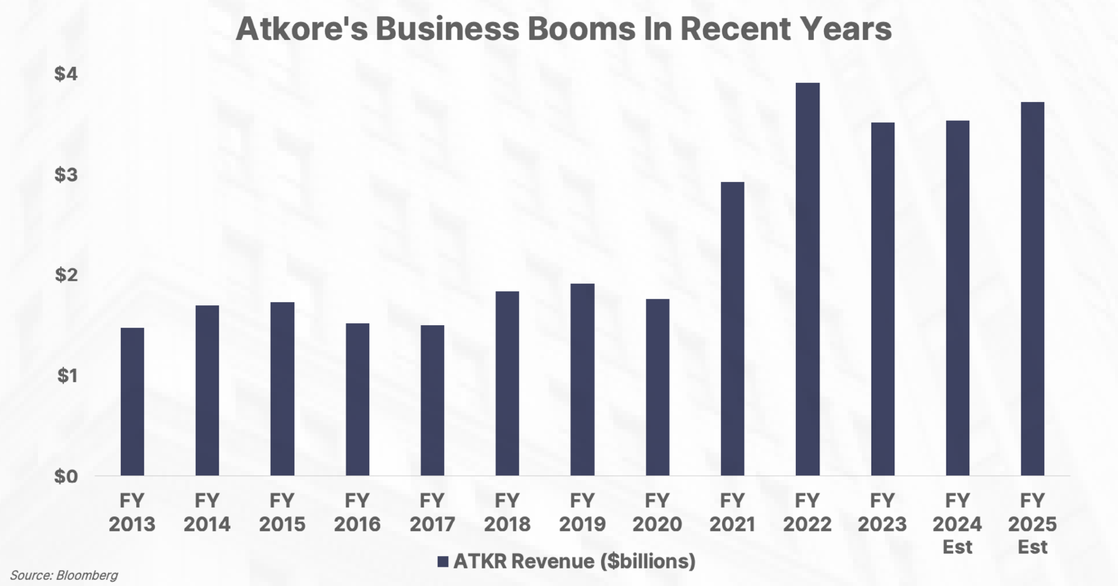

A retreat in the inflation-driven price increases for construction materials caused Atkore's 2023 revenue decline. When the flood of stimulus money began pouring into infrastructure spending from 2021 to 2022, it bumped into pandemic-era labor shortages and supply-chain disruptions. This triggered significant price inflation in construction costs that boosted Atkore's revenues. Then, as the labor market and supply chains began normalizing in 2023, prices retreated, causing a decline in both revenue and profit margins.

However, as we'll show later, Atkore has maintained substantial pricing power thanks to a powerful competitive advantage in its business model. But first, let's dive deeper into the key market trends that will fuel massive demand for Atkore's products, and continued revenue growth for at least the next decade.

The EV and AI Revolutions Provide a One-Two Punch

One major driver of electrical infrastructure demand will come from the electric-vehicle ("EV") revolution.

Despite our **bearish view** on the share price of the leading U.S. EV maker, Tesla, there's no denying that the broader EV revolution is here to stay. The federal government gave EVs a major push forward with the $1.2 trillion IRA legislation, which authorized a $7,500 tax credit for eligible EVs purchased through 2032.

Because of the generous government handouts and tax incentives at both the state and local level, financial services firm S&P Global expects EVs will make up 40% of new vehicle sales by 2030. That's a five-fold increase from the 8% market share EVs achieved in 2023. The firm also expects the total number of EVs on the road will reach 27 million, up 10-fold from around 2.5 million currently.

This increase in EVs means the number of electrical charging stations will also need to increase by a similar amount. According to analysis from consulting firm PricewaterhouseCoopers, the number of EV charging stations in the U.S. will need to grow from around 4 million currently to 35 million by 2030.

Demand for Atkore's electrical products has risen alongside the growing demand for EV chargers. And it's poised to continue expanding rapidly alongside the rising demands of EV charging, as management explained at a conference with investment analysts in 2023:

> *"It doesn't matter what EV charger. What I can guarantee you is there will be electrical lines, cable, conduit hooking up to that. So no matter what product gets installed, our products will be carried along with that, all the different infrastructure."*

Next, there's the growing demand for power generation from the AI revolution. A recent analysis from the International Energy Agency ("IEA") showed that the average ChatGPT AI query consumes nearly 10x the energy as a typical Google search. The agency expects power demand from the AI boom to increase at least

10-fold between 2023 and 2026 alone. Other experts have projected that data-center power demand will consume roughly 25% of all U.S. electricity over the next decade, *up from just 4% currently*.

The problem is, new data centers are popping up faster than new power plants. As a result, data-center construction is facing significant delays due to lack of electricity to feed their energy-hungry operations. Commercial real estate firm CBRE recently reported that project timelines for data centers have been extended by two to six years due to lack of power – in some cases doubling the construction timeline.

The Largest Electric-Grid Overhaul in Generations

In order to meet the ravenous power demands for the EV and AI revolutions, experts project that the U.S. will need to bulk up its electricity transmission capacity by 60% through 2030. This will require a massive undertaking, especially considering that U.S. power generation has remained flat for the last 15 years. As a result, the U.S. electric grid is on the verge of its biggest overhaul in generations.

In October 2023, the Biden administration announced the largest-ever investment into upgrading the country's electric grid. The IIJA and IRA legislation granted $30 billion to the Department of Energy to invest in 58 projects across 44 states, all dedicated to beefing up America's electrical grid.

But this is just a small drop in the massive bucket of investment that will be required in the coming years. The investment needs go beyond simply building more power stations and electrical wiring. The U.S. electric grid is old and antiquated. Nearly 70% of the U.S. power transmission infrastructure is over 25 years old. Another problem is that most power lines that move electricity through the grid were installed above ground. While this made for a cheaper upfront installation cost, it's turned into a massive long-term liability.

Consider the situation in California, where overgrown trees routinely knock over power lines and spark devastating wildfires. Pacific Gas & Electric (PG&E), one of the country's largest utilities, has been plagued by a series of destructive wildfires caused in part by trees falling into its air-suspended power lines. This includes the 2018 Paradise fires that killed 85 people and destroyed over 18,000 structures. PG&E filed for bankruptcy protection in 2019 after racking up $30 billion in damages for fires attributed to its electrical equipment.

The reorganized entity that emerged from bankruptcy protection is now proposing a massive $5.9 billion investment into burying 10,000 miles of power lines over the next decade. That's just one utility company. Similar efforts are underway across the country, including in Florida, where hurricane winds routinely destroy above-ground power lines – cutting off electricity for large swaths of the population.

This process of burying power lines, known as "grid hardening," will require a huge investment of between $200 billion and $300 billion over the next decade. These buried power lines will require the protective PVC conduit Atkore makes to keep underground wiring insulated and protected against the elements.

Atkore is the number-one supplier of PVC conduit in the U.S. Thus Atkore will become one of the major beneficiaries from the tens of thousands of miles of power lines installed underground over the next decade.

Demand for American-Made Solar Equipment Doubles Overnight

In Atkore's Safety & Infrastructure segment, one key product segment includes torque tubes, a component for solar power installations. These cylindrical metal tubes run along the axis of solar arrays, allowing the panels to tilt and rotate to maintain the proper angle with the sun. By providing the maximum exposure to sunlight, these torque tubes are critical for optimizing the efficiency of solar energy generation.

Because the Biden administration and environmental groups are advocating for a greater reliance on alternative energy, massive sums of money are pouring into this industry. Fueled by generous government subsidies, including a boost from the IRA legislation, solar-power installations will more than double over the next decade:

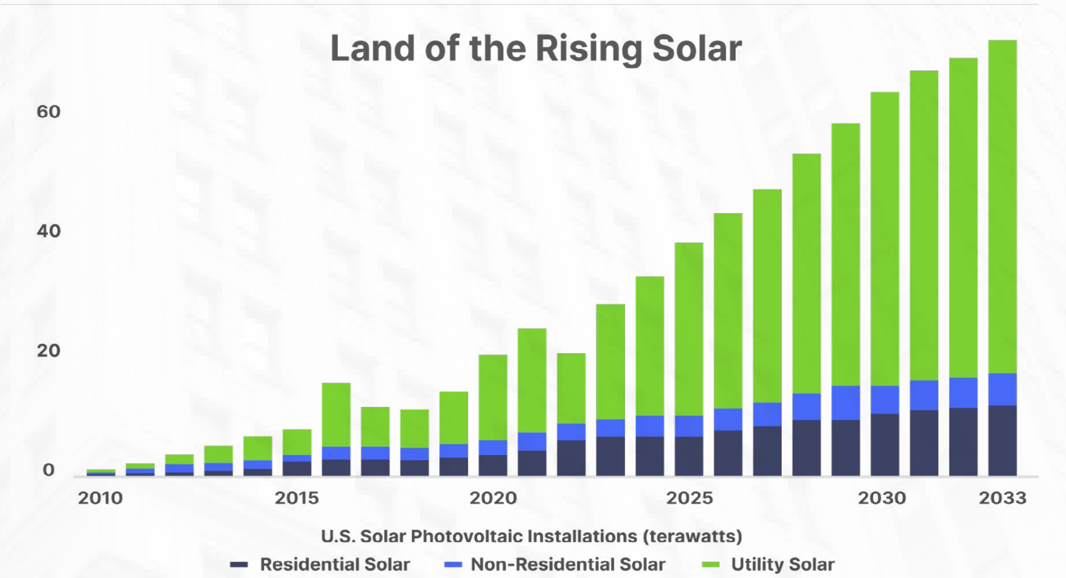

Source: Wood Mackenzie

But the growth in U.S. solar installations is only half the story fueling the demand for Atkore's torque tubes. As part of the IRA legislation, the Biden administration provided a series of tax and other regulatory incentives that heavily favor the domestic manufacturing of solar-energy components, including torque tubes.

Prior to these IRA incentives, a substantial portion of U.S. torque tubes were manufactured overseas, mostly in China. After the IRA legislation became law in August 2022, the economics simply don't make sense for overseas manufacturers.

As a result, the market for domestically produced torque tubes doubled overnight, as management explained in its Q1 2024 earnings call:

> *"With the Inflation Reduction Act, it should move all the [torque tube] volume into the states, which is a great thing for the U.S. and its economy... even if the solar market did not grow..., this doubles the amount of [domestic] solar torque tubes. And right now, I don't think that capacity exists by anybody out there."*

With a doubling in U.S. solar-power installations over the next decade, plus a doubling in the market share for domestic suppliers, Atkore's torque-tube business is booming. And it was perfectly positioned for this growth spurt. A year before the IRA legislation passed in August 2022, Atkore was already putting money to work expanding its torque-tube capacity. In May 2022, it expanded its Phoenix, Arizona, manufacturing facility to add two new production lines for solar torque tubes. The company is also investing in its Hobart, Indiana, plant to produce more solar torque tubes and other metal tubing products.

The Broadband Boost

Another key demand driver for Atkore's is in the flexible conduit made from a polymer plastic known as High Density Polyethylene ("HDPE"). HDPE is used primarily as a protective material for fiber-optic cable used in telecommunications, internet, and data-center applications.

From December 2021 to November 2022, Atkore acquired four leading U.S. manufacturers of HDPE conduit, plus another company specializing in HDPE recycling. Atkore went from zero to the second-largest HDPE conduit supplier in the U.S. within just two years.

As it made these acquisitions, HDPE conduit demand was already growing rapidly, thanks to the expansion in 5G wireless internet and the AI-created boom in data-center demand. But the market received another massive boost in June 2023 when the Biden administration announced $42 billion in federal funding for high-speed internet expansion throughout the U.S. as part of the IIJA legislation.

This will be the largest-ever investment into expanding high-speed internet access – with the government aiming to bring reliable broadband to an additional 8.5 million households and small businesses. All of these new internet connections will require fiber-optic cable, and those cables will require protection in the form of HDPE conduit. As the second largest HDPE conduit supplier in the U.S., Atkore will benefit from this new source of demand.

The total market size in HDPE conduit today is roughly $7 billion, a figure that Atkore estimates will double by 2029.

Across all of Atkore's product segments, the company estimates its total addressable market size is around $40 billion. And given the massive long-term growth trends across virtually every major business line, this will no doubt explode in size over the next decade and beyond.

Atkore's Trifecta: Growth + Margin Expansion + Share Count Reduction

From an investor perspective, growth is only a small part of the Atkore opportunity. The potential returns from this business go well beyond the rate of revenue increases.

Consider the following...

Since Atkore's 2016 IPO through the end of 2023, revenue grew 133% from $1.5 billion to $3.5 billion, about 13% annually. A decent result, but nothing extraordinary. However, over that same time period, the company's share price has delivered a massive 1,092% total return. This compares with a 186% gain in the S&P 500 over the same period.

The key factor behind the remarkable rise in the share price is the fact that the business has become increasingly more profitable and more capital efficient since 2016. Profit margins have increased nearly five-fold, from 4% in 2016 to 19% today. Free cash flow ("FCF") margins have nearly doubled from 9% to 17%. Atkore has returned that growing FCF into share buybacks, reducing its share count by 40%, from 62 million in 2016 to 37 million currently.

The trifecta of revenue growth, margin expansion, and a falling share count has boosted Atkore's earnings per share by nearly 20-fold, from $0.94 in 2016 to $17.50 in 2023.

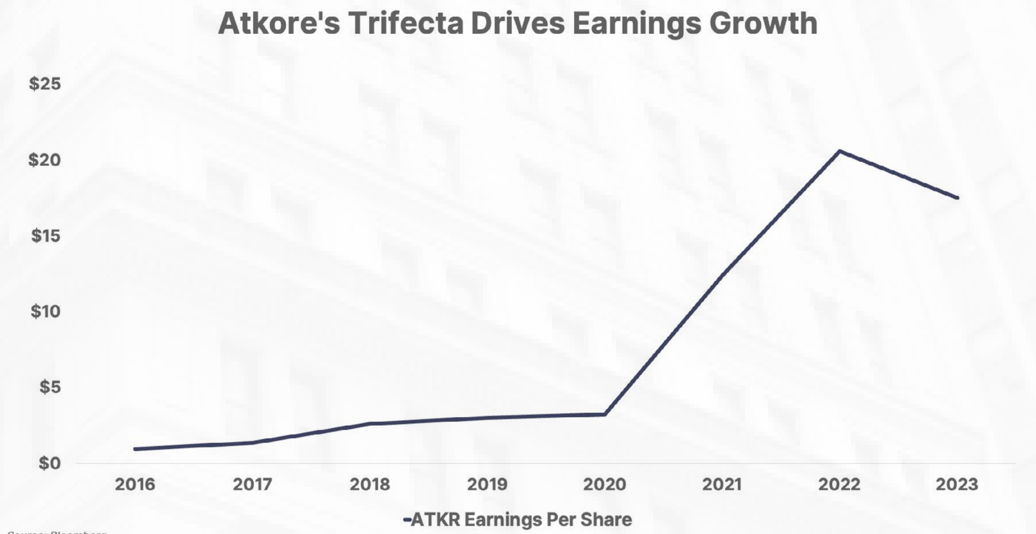

Source: Bloomberg

Now here's where things get interesting. Despite this incredible improvement in the business fundamentals, the market has given little respect to Atkore. The shares currently trade at around 10x earnings, or a 50% discount to the S&P 500's 20x earnings multiple.

Meanwhile, speculators have accumulated a massive short position in Atkore's shares, equal to nearly 14% of the shares outstanding:

Source: Bloomberg

The speculators' view, as far as we can tell, is that Atkore's expanding profit margins since 2019 were a one-off result of pandemic-driven supply constraints. Going forward, bears expect that Atkore's profit margins will return to pre-pandemic levels now that supply chains have normalized. In this scenario, Atkore's earnings could suffer a 65% drop as profit margins fall from 19% now to 2019 levels of around 7%.

However, as we'll show in the next section, we believe the bears are mistaken and that Atkore's industry-leading profit margins are here to stay.

The Danaher of Electrification

The secret to Atkore's growing profitability lies in a set of operating principles known as the Atkore Business System ("ABS"). That name might ring a bell, as it bears a striking resemblance to the Danaher Business System ("DBS"), which we wrote about in our **March 22 recommendation of buyout firm Danaher**. That's no coincidence. The man who developed the ABS, John Williamson, spent over a decade at Danaher before becoming Atkore's CEO from 2011 to 2018.

ABS incorporates many of the same principles that made Danaher one of the greatest wealth compounders of all time. This includes using the *kaizen* manufacturing principle of making continuous operational improvements to minimize waste, boost profit margins, and improve customer service.

But Atkore's biggest competitive advantage has come from replicating Danaher's acquisition playbook. Specifically, following the strategy of acquiring a group of the leading brands in the same or similar industry segments, as we described in our Danaher recommendation.

Since going public, Atkore has made 18 acquisitions. This includes bulking up business lines it already participated in, like PVC conduit. Through four acquisitions since 2016, Atkore has become the number-one market-share leader by a wide margin. It now has 10 PVC manufacturing plants, compared to just four for its next closest rival. This gave Atkore an unmatched distribution footprint, with a PVC plant within 500 miles of every customer across the country.

Likewise, when Atkore entered the HDPE conduit market with its first acquisition in December 2021, it quickly scooped up three additional manufacturers – taking it to the number-two market-share position within just two years.

Across all major business lines, Atkore now holds the number-one or two market-share position.

This growing scale creates several key advantages. First, it allows Atkore to buy materials in larger quantities, and thus reducing costs. Second, the growing concentration of facilities reduces the distance between Atkore and its customers. Since many Atkore products are large, bulky items, this reduces transportation costs.

But reducing costs through economies of scale is fairly commonplace – all of Atkores competitors do the same. What sets Atkore apart is that it has found a way to lower its customers' operating costs. This, in turn, has allowed it to charge premium prices and generate sustainability higher profit margins than the competition.

In order to understand how this works, we must first understand the ecosystem of competitors and customers that Atkore operates within.

Gaining Pricing Power in a Commodity Business

The vast majority of Atkore's sales (83%) are made to large-scale electrical-parts distributors. These distributors buy products from manufacturers like Atkore, and then resell them to the end customers for use in construction projects.

Most of these distributors use what's known as a hub-and-spoke distribution system. In this system, distributors stockpile inventory in large warehouse hubs, then feed products through a series of smaller distribution warehouses. From there, distributors sell products to the end customer, typically construction firms and contractors.

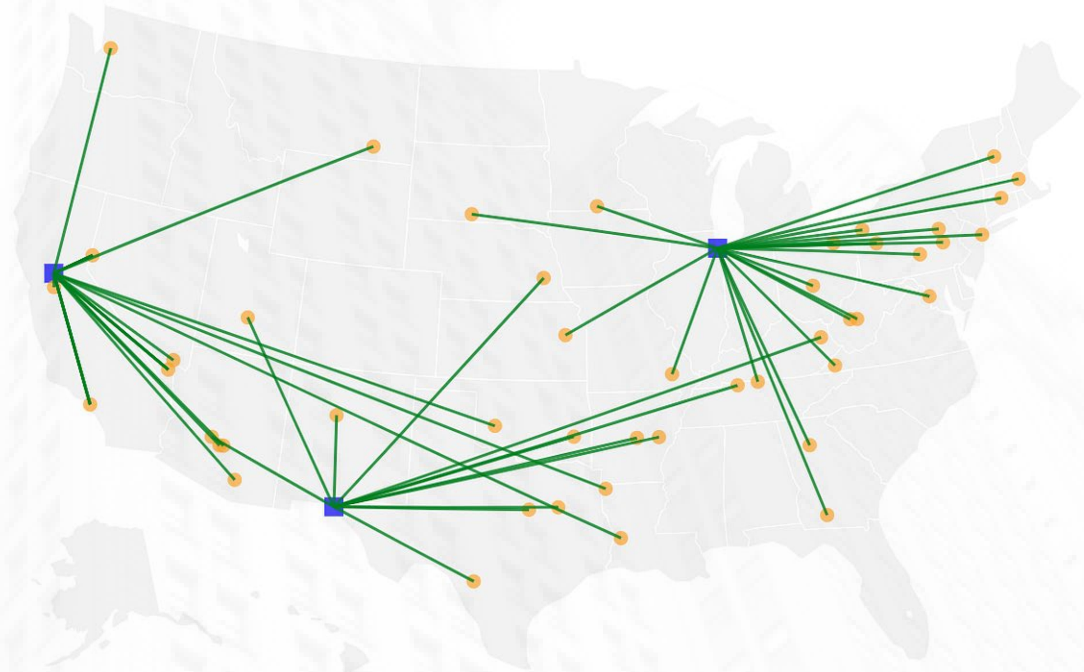

Atkore separates itself from its competitors in that it can bypass sending products to the hubs of its customers' distribution centers, and instead ships products directly into the smaller warehouse locations. It achieves this through what it calls a "one order, one delivery, one invoice" service, which packages multiple products – up to seven at a time – onto a single truckload for delivery.

This bundling feature allows Atkore to economically ship smaller quantities of materials, instead of sending a full truckload of a single product. In this way, Atkore still retains the economic benefit of transporting full truckloads of materials. But it enables its customers to receive those items in the smaller quantities needed to avoid overwhelming their smaller distribution warehouses.

Conversely, Atkore's top competitors typically ship full truckloads of single products at a time, because they don't offer the same breadth of product offerings. Thus, these competitors are forced to send these large, single product shipments into the hubs of their customers' distribution systems.

Shipping directly to the smaller facilities is much more valuable for distributors, because it eliminates their cost of moving products from hubs to their smaller warehouses. This also allows distributors to avoid tying up their precious capital into stockpiles of excess inventory that sit unsold at their distribution hubs for weeks or months at a time. By delivering products to a higher value location for its customers, Atkore has achieved five to seven percentage points of additional pricing power for the same products it used to sell in bulk on a single truckload.

And that's the big secret about Atkore: its distribution advantage unlocks premium pricing power from a commodity product. It's the most powerful feature of its business model, and one that most investors don't fully appreciate. As Atkore CEO Bill Waltz explains:

> *"And I do think if there is one thing that's misunderstood with Atkore from an investor perspective, it is that 'one order, one delivery, one invoice.'... Our products are so big and bulky that these small distribution locations don't have the space or are really space constrained to bring a truckload of steel conduit or a truck load of PVC. So where we can combine all those products together on one truckload, it reduces their inventory, reduces the freight expense, and reduces the space they need for those products."*

Atkore introduced its order bundling feature in 2019. Since then, its profit margins soared from 7% to 23% during the peak of inflation-driven pricing gains in 2022, before falling to a still-impressive 19% in 2023.

This brings us back to the short-sellers' bearish argument on Atkore – that its recent gains in profit margins came from a one-time pandemic-driven surge that will dissipate. The truth is, the inflationary outbreak in construction materials has already passed. After reaching highs of over 30% inflation per year in 2021 and 2022, prices for construction materials entered into deflationary territory in November 2022, and have remained roughly flat since then. And yet, Atkore has maintained its profit margins at permanently higher levels.

Source: U.S. Bureau of Labor Statistics

For a final perspective on profit margins, let's consider Atkore versus its top competitors. All of the construction-supply companies enjoyed the same pandemic-driven surge in pricing and profitability. So by comparing Atkore's profit margins versus its top competitors from 2021 to 2023, we can eliminate any distortions from macro trends. When we do this, it's clear that Atkore enjoys a meaningfully higher margin advantage versus its key competitors:

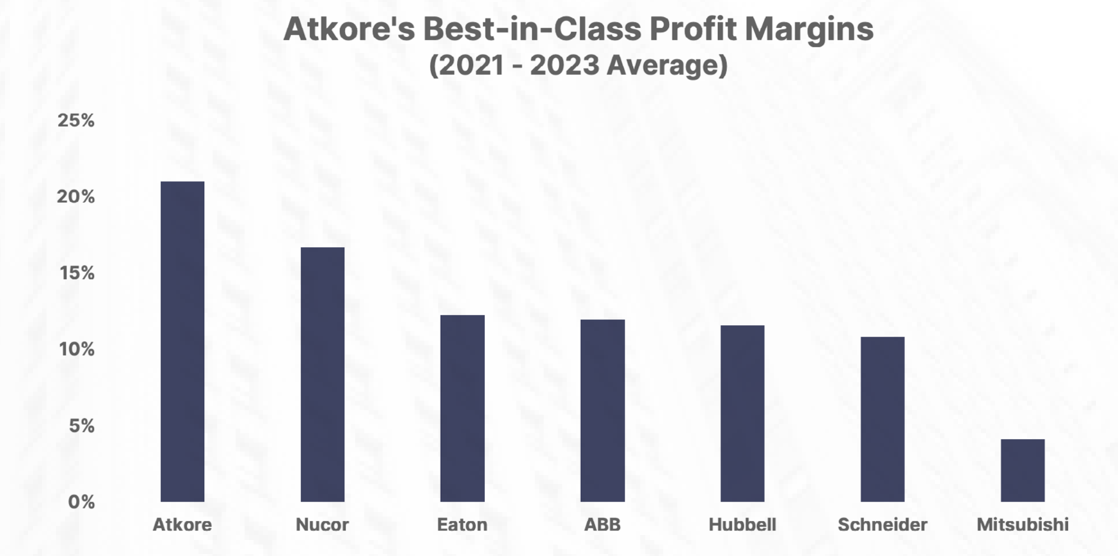

Source: Bloomberg

Looking ahead, all signs indicate Atkore's industry-leading profit margins are here to stay.

An Enduring and Growing Competitive Advantage

The reason Atkore will maintain its competitive advantage lies in the differentiated approach it's taken in organizing its business. Atkore has strategically acquired a portfolio of products designed to cater specifically to the major distributors of electrical construction products.

That is, the company has pursued what's known as horizontal integration – acquiring a collection of businesses designed to sell different products into the same end market. Conversely, Atkore's major competitors have all pursued vertical integration – they've built their businesses around fully integrating various stages of the production process into making the same product.

Consider steel producer Nucor, which is Atkore's largest competitor selling metal framing, cable, and conduit products to electrical equipment distributors. Nucor has focused on vertically integrating every step in the steel-making process, turning the raw minerals of iron ore, limestone, and coal into finished products. This vertical integration makes Nucor the largest and lowest-cost steel manufacturer in America.

This means that Nucor can produce things like steel framing and conduit cheaper than Atkore can. But it can't deliver the same value that Atkore delivers to its customers, because it can't bundle multiple products and get these products to their highest-value location as efficiently as Atkore.

The same can be said of Atkore's other key competitors that are narrowly focused on PVC or HDPE conduit and electrical-cable wiring. Each of these businesses has invested around their individual products, versus Atkore's approach of investing around the needs of its key customers. And so far, none of Atkore's competitors have shown any desire to replicate its horizontally integrated strategy, as management has explained:

> *"The difference between Atkore and our competitors is that we are built around supplying the electrical industry. If you think of somebody like Nucor – a great company, but they're vertically integrated... I just can't imagine the general manager at Nucor going to the board of directors and going, 'hey, I had this opportunity to get into plastics.'...So, it's not even in their mindset."*

Even if Atkore's competitors changed their strategy in an attempt to replicate Atkore's business model, they would be starting from a significant disadvantage. Atkore has already consolidated the majority of markets it operates in, becoming number one or number two across its largest product segments. And it's continuing to invest aggressively in furthering its lead – as seen in Atkore's aggressive entrance into the HDPE market, where it rose to the number-two market-share leader in the span of just two years.

Finally, because Atkore has already established the dominant horizontally-integrated business model for electrical infrastructures supplies, it can afford to pay a premium when making an acquisition that its competitors can't justify. That's because Atkore can immediately incorporate acquired businesses into its bundled product offerings, and instantly realize pricing power over its non-horizontally integrated competitors.

From a shareholder perspective, Atkore's acquisitions are adding increasingly more value to its overall business by further expanding its distribution dominance and economies of scale. We can see this in Atkore's improving return on invested capital over time, which has more than tripled from 10% in 2016 to 32% in 2023:

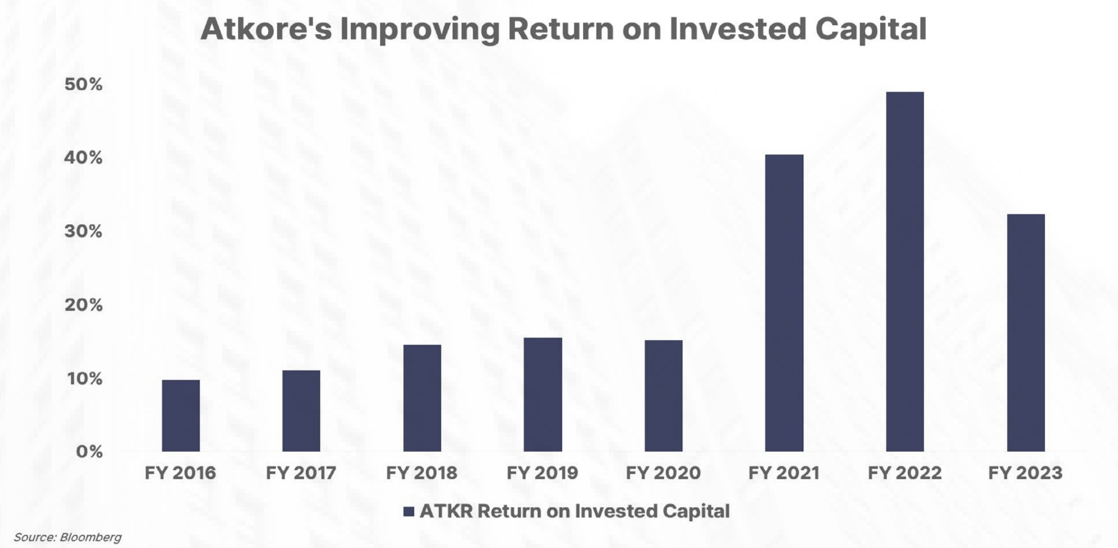

As a leading provider of the critical materials that will fuel this electrical-infrastructure boom, Atkore is perfectly positioned to benefit from this powerful trend.

We hope you found this report interesting and compelling. Subscribers to *The Big Secret on Wall Street* (and our Partner Pass members) have access to our latest investment research, including the securities discussed in this report. To learn more, call Lance James, our Director of Customer Care, at **888-610-8895**, internationally at **+1 443-815-4447**, or click the link below.

Click here to sign up for *The Big Secret on Wall Street* →