

Three Chapters on Where the Returns Actually Come From

A sampler from the investing library of Victor Wong

About this sampler

Three chapters, from three different books, chosen because they stand on their own.

The first asks why so few stocks now decide what a whole market does. The second is the oldest idea in investing — buying a business for less than it is worth — and why it still works when the business is a semiconductor company. The third is the practical one: how to actually buy the AI super-cycle without overpaying for it.

Read them in order, or start with whichever question is yours. They are complete chapters, not extracts.

1. Why So Few Stocks Matter

“Diversification is protection against ignorance. It makes little sense if you know what you are doing.”

— Warren Buffett, Berkshire Hathaway annual meeting, 1996

At first, it looked like a market rally. Then the numbers became too large to ignore.

Nvidia reported **\$27.0 billion** in revenue for fiscal 2023. One year later, revenue had surged **126%** to **\$60.9 billion**. The following year it more than doubled again, reaching **\$130.5 billion**. In two fiscal years, annual revenue grew almost fivefold.

This was not a normal corporate success story. It was a signal flare. The market understood that something enormous had changed. What it got wrong was where the change was happening.

Almost fivefold

Revenue by fiscal year,

\$27.



Figure 1-1. Revenue almost quintupled in two fiscal years, and the fiscal 2025 step added more dollars than the company’s entire fiscal 2024 revenue.

Investors began asking who supplied the engines, the fuel systems and the roads of the new industrial platform, and the answer pointed to a surprisingly small group of companies.

The Numbers That Changed the Conversation

Table 1-1. The five companies and the figure each one highlighted.

Company	The figure	Why it mattered
Nvidia	Revenue rose from \$27.0B in fiscal 2023 to \$60.9B in fiscal 2024, then \$130.5B in fiscal 2025.	AI computing demand was becoming visible in company-sized leaps.
SK hynix	2023 revenue ₩32.77T with an operating loss of ₩7.73T ; 2024 revenue ₩66.2T , operating profit ₩23.5T ; 2025 revenue ₩97.15T , operating profit ₩47.2063T — a 49% operating margin.	Advanced AI memory was a strategic bottleneck.
Samsung Electronics	Operating profit climbed from ₩6.57T in 2023 to ₩32.7T in 2024—nearly five times higher—then ₩43.6T in 2025 on revenue of ₩333.6T .	The memory recovery could alter the earnings power of a giant in a single year, and then keep going.
TSMC	2024 revenue rose 33.9% and net income rose 39.9% .	The factories at the heart of advanced chips captured more of the AI build-out.
Broadcom	Fiscal-2024 AI revenue grew 220% year over year.	The boom extended beyond processors into the networks that connect AI systems.

The sc

One-year

Samsun

Broadc

Nvidia

TSMC r

Figure 1-2. Five measures of one year's growth, each named as itself — a rise in operating profit is not the same claim as a rise in revenue.

These are company results, not promises about future share prices. They explain why the stocks became so important: financial markets are always trying to price the future, and when sales, profits and investment plans begin moving at this speed, the future stops feeling distant.

The Conclusion the Market Drew

The first conclusion investors drew was the wrong one, and it was wrong in an instructive way. If artificial intelligence was a software revolution, then the companies that profited most should look like software companies: lightweight, fast-moving, and able to serve a new customer at almost no additional cost. That is how the internet's first decade worked, and it is the model most investors had in mind.

The companies that actually moved were nothing like that. Nvidia designs chips and depends on factories it does not own. SK hynix and Samsung build memory in plants that cost many billions and take years to bring online. TSMC turns designs into silicon inside buildings full of machinery that must be replaced on a schedule set by physics rather than by fashion. These are industrial businesses with industrial risks: construction schedules, equipment lead times, manufacturing yields, and the price of electricity.

That difference explains why the numbers grew so violently, and why the effects travelled so far. A software business can double its revenue without doubling its costs. An infrastructure business cannot, because every extra unit of demand has to be met with concrete, copper, water and machinery, none of which can be conjured inside a quarter. When demand arrives faster than capacity, prices rise, and the suppliers of the scarcest parts of the chain collect most of the difference.

There is a second consequence, and it is the one that reaches most readers. Because the profits concentrated in a handful of companies, the indices concentrated with them. An investor who bought a broad American index fund in 2020 was buying a diversified basket of businesses. An investor who bought the same fund in 2026 was buying something much closer to a bet on a single industry. The fund did not change. Its contents did.

None of this required anyone to forecast the future. It required reading the results the companies published, and noticing how few companies were doing the spending.

From a Screen to a Supply Chain

It is tempting to think of AI as something that lives on a screen. But every answer generated by an AI system has a physical history. Before a model can write a sentence, a chip must perform calculations; before that chip runs at full speed, memory must feed it data; before the memory can be made, chemical materials, precision machinery and engineering expertise must arrive at a factory; and before the factory can expand, someone has to decide that demand will last long enough to justify spending billions.

That is why a few stocks command so much attention. They are claims on the narrowest passages in the global technology supply chain.

Every answer has

One layer of the stack, and

1

Com
Nvid

2

Hig
SK H

Figure 1-3. The chain runs in one direction, and a few companies sit on its narrowest passages.

Nvidia — the compute layer. Nvidia sits on the narrowest passage of all. Its processors are unusually good at one thing: performing enormous numbers of small calculations at the same time, which happens to be exactly what training and running a model requires. The company did not arrive there by accident. It had also spent nearly two decades building **CUDA**, the software layer that lets programmers use those processors for something other than graphics, and that software is why a faster chip from a rival still struggles to displace it. Its revenue went from **\$27.0 billion** to **\$60.9 billion** to **\$130.5 billion** in two fiscal years. That is what it looks like when the world's largest industry decides it needs one company's product faster than that company can make it.

SK hynix — the fuel system. A processor with nothing to read from is an expensive heater. Training a model means moving enormous volumes of data between the processor and its memory, and if the memory is too slow the processor spends its life waiting. High-bandwidth memory solves that by stacking memory dies and wiring them directly to the processor, putting the memory maker inside the processor's own package. It is not a component a data-centre operator can quietly swap for a slower one; substituting it means giving up much of the performance that was paid for. SK hynix led this market for years, and in 2025 it earned an operating profit of **₩47.2 trillion** — more than the entire Samsung group, on less than a third of its revenue.

Samsung — the crossroads. Samsung sits in two places at once, making memory and manufacturing chips for others, and its numbers show how violent that position can be. Operating profit fell to **₩6.57 trillion** in 2023, recovered to **₩32.7 trillion** in 2024, and reached **₩43.6 trillion** in 2025 on revenue of **₩333.6 trillion**. For a company that size, those are not ordinary swings. They are what happens when a supplier is essential to a build-out whose demand arrives in steps far larger than any factory can be built in.

TSMC — the factory everyone needs. TSMC does not design the chips. It makes them, for almost everyone who does, and it makes more than **72%** of the world's outsourced semiconductor production. A leading-edge logic plant is among the most expensive objects a private company can build and is worth nothing until it is finished — which is why the position has proved so durable. An order is not a plant, and no amount of money buys back a lost year. In 2024 its revenue rose **33.9%** and its net income **39.9%**. Every layer above it — the designs, the accelerators, the systems built around them — passes through its buildings.

Broadcom — the wiring. Broadcom is the quiet one, and the most instructive. It does not sell an accelerator. It sells the parts that make accelerators useful together: the networking that moves data between thousands of them, and custom chips built to a customer's own specification. Its AI revenue grew **220%** in fiscal 2024 — a faster rate than the company whose chips it connects. That is the pattern the rest of this book keeps returning to: the shortage moves, and the money follows it to whichever layer of the chain is scarcest at that moment.

Trough, recove

Operating profit in trill

~~¥~~6.57T




Figure 1-4. Three years, two companies. Samsung went 6.57, then 32.7, then 43.6 trillion won. SK hynix went from a loss of 7.73 trillion to 23.5 and then 47.2 — overtaking the whole of Samsung in 2025, on ₩97.1 trillion of revenue against Samsung's ₩333.6 trillion.

Who Competes at Each Layer

Naming the leader at each layer invites the obvious question, and it decides whether these are permanent monopolies or merely profitable temporary ones: who else can do this?

Table 1-2. Each layer, its leader, and who is trying to take it.

Layer	The leader	The challengers	How close they are
Compute	Nvidia, about 81% of AI data-centre chips (IDC, 2026)	AMD, Intel, and the customers' own chips	AMD's data-centre revenue was \$5.8bn in the March 2026 quarter, up 57% — real, and still a fraction of Nvidia's \$75.2bn in the same quarter. The faster-growing threat is not AMD at all: custom chips grew 44.6% in 2026 against 16.1% for ordinary GPUs
High-bandwidth memory	SK hynix, 64% → 58% → 50% in three quarters	Samsung 15% → 21% → 33% , Micron, CXMT	The gap narrowed from 49 points to 17 in a year — the clearest case of a chokepoint being broken
Memory and manufacturing	Samsung	SK hynix, Micron, Kioxia, CXMT	Samsung leads on size and lost the profit race: SK hynix earned ₩47.2tn in 2025 on a third of the revenue
Leading-edge foundry	TSMC, 72.5% of the market (TrendForce, Q2 2026)	Samsung Foundry 5.9% , SMIC 5.4% , UMC, GlobalFoundries	The gap is widening , not closing: TSMC and Samsung were 66.6 points apart, up from 65.8 the quarter before. Intel does not appear in the top ten
Networking and custom chips	Broadcom, AI revenue heading for about \$56bn in fiscal 2026	Marvell, Nvidia's own networking, Cisco and Arista	Broadcom went from \$20bn to a guided \$56bn in one year — the fastest-growing position in the chain

Now the same five layers in plain terms.

Compute: the challenger is real, and it is not the one you would expect. AMD has done what almost nobody else has managed: built a second source for AI processors that large customers actually buy. Its data-centre revenue rose **57%** in the March 2026 quarter, Meta and OpenAI have both taken warrants tied to buying its chips, and it now holds something between **7% and 10%** of the market depending on how you draw the boundary. But look at the absolute numbers rather than the percentages. In the same quarter, Nvidia's data-centre revenue was **\$75.2 billion** against AMD's **\$5.8 billion**. AMD is winning the right to exist as a second source while the gap in dollars widens. Intel, which once supplied the processors inside almost every server in the world, does not register in AI accelerators at all — and in January 2026 it told investors it could not keep up with demand even for its own server chips, leaving profitable business on the table. The genuinely dangerous competitor is a different kind of company: the customers. Google, Amazon, Microsoft and Meta are all designing their own accelerators, and it is Broadcom and Marvell who build them. **Custom chips grew 44.6% in 2026; ordinary GPUs grew 16.1%.** The threat to the leader is not a rival product. It is the buyer deciding to stop buying.

Memory: a chokepoint that visibly broke. This is the most instructive layer in the book, because the change happened in public and in under a year. SK hynix held **64%** of high-bandwidth memory; it now holds about **50%**, while Samsung went from **15%** to **33%** and Micron, which lost money on memory two years earlier, reached a gross margin of **84.6%**. A shortage that looked structural turned out to be temporary, and the market repriced all three companies as the shares moved. Chapter 6 follows this race quarter by quarter.

Samsung's position: leader by size, loser by margin. Samsung makes more memory than anyone and manufactures chips for other people as well, and both halves of that position are under attack at once — from SK hynix in memory and from TSMC in manufacturing. Its own numbers show the strain: record operating profit of **₩43.6 trillion**, and simultaneously a foundry business with **5.9%** of the market and falling.

Foundry: the gap is widening, which is the opposite of what competition is supposed to do. Samsung Foundry is profitable on its older 4-nanometre lines and improving its 2-nanometre yields, and none of it is closing the distance. TSMC's share went **up** while Samsung's went **down**, and the two are now **66.6 percentage points** apart. SMIC, China's largest, grew 20% in the quarter and still holds only **5.4%**. Intel, with the largest capital programme in the industry behind it, has not yet broken into the top ten at all. This is what a maturing monopoly looks like: challengers spending heavily, improving, and losing ground.

The foundry market

Share of pure-play foundry

— TSMC — Samsung

80%

60%

40%



Figure 1-5. The foundry market over six quarters. TSMC's share rose while Samsung's and SMIC's fell, and the gap between first and second widened to 66.6 percentage points. Counterpoint Research, on a pure-play basis; TrendForce's broader all-foundry basis gives TSMC 72.5% and Samsung 5.9% for the same quarter, which is the pair quoted in Table 1-2.

Networking and custom chips: the fastest-growing position, and the least visible. Broadcom designs the parts that let thousands of accelerators work as one machine, and it designs custom chips for customers who want to stop buying a standard one. Its AI revenue is heading for about **\$56 billion** in fiscal 2026, against roughly **\$20 billion** a year earlier. It does not appear in most books about artificial intelligence, and it is the clearest example of the pattern this chapter keeps returning to: when one layer becomes scarce, capital moves to the next one.

The pattern across all five is worth stating plainly. In four of the five layers the leader is holding or extending its position. In the fifth — high-bandwidth memory — the leader lost fourteen points of share in a year, which proves the others can be broken. That is why the rest of this book follows the chain rather than the companies: what matters is not who is on top now, but which layer is about to become the scarce one.

The Pattern Before

This is not the first time a real technological transformation has rearranged the world's stock markets. Railways did it in the nineteenth century, electrification in the early twentieth, and the fibre-optic buildout of the late 1990s in living memory. In each case the technology delivered on its promise, and in each case a large number of the investors who funded it did not make money. An industry being right is not the same thing as a stock being right, and the difference is usually the price that was paid.

So the honest position at the start is neither enthusiasm nor alarm. The transformation is real, the spending is enormous, and the concentration it has produced in the world's indices can be measured rather than argued about. What remains open is whether the returns will justify the prices, and that is a question this book follows through the supply chain rather than settles in advance.

A Rally With Consequences

The phrase “a few stocks move the world” sounds exaggerated until you follow the consequences. When Nvidia's results surprise the market, its suppliers and competitors move with it. When TSMC reports stronger demand for advanced nodes, the message travels to chip designers, cloud companies and policymakers on the other side of the Pacific. The share price is only the visible spark. Behind it stands a chain of expectations: more data centres, more electricity, more factory construction, more advanced packaging, more trained engineers, and more competition between countries that do not want to depend on someone else for the technologies of the future.

Markets can price a technological shift too early, and the history of semiconductors is full of booms

followed by painful corrections. But a correction would not make the underlying question disappear. The world is building new infrastructure for artificial intelligence, and a few companies control some of its most difficult components. Their stocks matter because the systems behind those stocks matter.

2. The Economic Moat

"In business, I look for economic castles protected by unbreachable moats." — Warren Buffett, 1995 Berkshire Annual Meeting

The Miracle of Omaha

There is a city in the American Midwest called Omaha, Nebraska. It is not a financial capital. There are no skyscrapers full of investment banks, no Wall Street trading floors, no hedge fund complexes. It is a quiet city where families have lived for generations, where the most exciting event of the year is the College World Series, and where the biggest attraction is a modest office building at 3555 Farnam Street.

From that building, for over sixty years, something impossible happened.

Between 1965 and 2025, a single man — working in a modest office with no computer, no Bloomberg terminal, no algorithm, no team of analysts — turned a failing textile mill into one of the largest companies in the world, worth over **\$800 billion**. He did it without leverage. He did it without insider trading. He did it without a single hostile takeover in his later years. He did it with fewer than twenty stocks at any given time, most of which he held for decades. And he did it at an annualized return of roughly **19.9%** — nearly double the stock market's 10.4% over the exact same period.

That man is Warren Buffett. The company is Berkshire Hathaway. And financial historians call the numbers **the Miracle of Omaha**.

To understand how extraordinary this is, let me walk you through what 19.9% compounded annually actually means.

In 1965, if you had invested **\$10,000** in the S&P 500 index — the broadest measure of the American stock market — and never touched it, reinvesting every dividend along the way, that money would have grown to approximately **\$1.2 million** by 2025. That is a life-changing sum. It would pay for a child's education, buy a home, or fund a comfortable retirement.

But if that same \$10,000 had been invested in Berkshire Hathaway stock in 1965, it would have grown to approximately **\$550 million**.

\$10,000 Invested in 1965 → 60 Years Later

S&P 500 Index → ~\$1,200,000

Berkshire Hathaway → ~\$550,000,000

*Berkshire outperformed the market by more than 450 times***

One man. One small office. No computer. No algorithm. Two principles that never bent. That is the miracle — not luck, not magic, but the relentless application of a few simple ideas over a very long time.

What were those ideas? The first — and the subject of this chapter — is the **Economic Moat**: what makes a business worth owning in the first place.

The Castle and the Moat

In 1972, Berkshire Hathaway bought a small California candy company called See's Candies for \$25 million. At the time, it seemed like an unremarkable transaction. See's was a regional chocolate maker with a few dozen stores, mostly in shopping malls. It had no technology, no growth story, and no Wall Street analyst covering it.

Fifty years later, See's had generated more than **\$2 billion** in cumulative profits for Berkshire. It had never required a single dollar of additional capital. It had survived recessions, wars, changing consumer tastes, and the rise of Amazon. It was never disrupted by technology. It simply kept selling boxes of chocolate, year after year, at higher prices, to customers who loved the brand.

Buffett called this the **economic moat** — a durable competitive advantage that protects a business from its competitors, the way a medieval moat protected a castle from invading armies.

"The most important thing to me is figuring out how big the moat is around the business. What I love, of course, is a big castle and a big moat with piranhas and crocodiles."

See's Candy had a moat. The brand was loved. The product was high-quality. Customers would walk past five other candy stores to buy a box of See's. Competitors could not replicate the taste, the reputation, or the customer loyalty. The moat was wide enough that no invader could cross it.

This chapter is about learning to recognize economic moats — and applying that skill to AI stocks.

The Five Types of Moats

Over decades of investing, Buffett and Munger identified five fundamental types of economic moats.

Understanding them is the difference between owning a business and gambling on a stock.

Moat Type	What It Means	Buffett's Best Example
Cost Advantage	You can produce the same product cheaper than anyone else	GEICO (lower distribution costs than traditional insurers)
Brand / Intangible Assets	Customers trust your brand and will pay more for it	Coca-Cola (brand worth \$100B+)
Switching Costs	Once customers use your product, it is too painful to leave	Apple (leaving the iPhone ecosystem means losing apps, data, and services)
Network Effects	The more people use your product, the more valuable it becomes	American Express (more merchants → more cardholders → more merchants)
Scale / Efficiency	You operate at a scale that smaller competitors cannot match	Burlington Northern (cost of replicating rail network is prohibitive)

Each moat type is different. Each requires different competitive dynamics. And each applies differently to AI companies.

Moat #1: Cost Advantage — The GEICO Story

In 1951, a 20-year-old Warren Buffett visited GEICO's headquarters in Washington, D.C. He knocked on the door on a Saturday, found a janitor who let him in, and spent hours talking to a GEICO executive named Lorimer Davidson.

What he learned changed his life. GEICO had a revolutionary business model: it sold car insurance directly to customers by mail, bypassing expensive insurance agents. This meant it could charge **30% to 50% less** than competitors and still make a profit. The cost advantage was structural — it was built into the business model, not dependent on any single manager.

Buffett bought his first GEICO stock in 1951 at age 21. But the real story — the one that reveals the power of a cost-advantage moat — happened 25 years later.

Here is the stock price journey that every investor should study:

Year	GEICO Stock Price	Event
1948	~\$27	Benjamin Graham's partnership buys 50% of GEICO; stock goes public
1951	~\$20	21-year-old Buffett buys \$10,000 worth (half his net worth)
1972	\$61	Peak of the bull market — the cost advantage was working perfectly
1974	\$42	First signs of trouble as rapid expansion strains reserves
1975	\$7	Losses mount; stock collapses as underwriting mistakes surface
1976 (Spring)	\$2	Near bankruptcy. Stock loses 97% from its peak. Buffett begins buying heavily
1977–1979	~\$5–\$10	Turnaround management takes hold; Buffett increases his stake
1980s	~\$20–\$40	Steady recovery as the cost-advantage model proves resilient
1995	~\$80	Berkshire buys remaining 49% for \$2.3 billion — a 40x return from the 1976 low
1996	N/A (taken private)	GEICO becomes a wholly owned Berkshire subsidiary

The most important row is **1976**. When GEICO's stock fell from \$61 to \$2, most investors saw a company about to die. Buffett saw a company whose **structural cost advantage** was still intact — the management had simply grown too fast and mismanaged the underwriting. The moat was still there. The problem was fixable.

He was right. The \$2 shares he bought in 1976 were worth over \$80 by 1995 — a **40-fold return** in less than 20 years. By 2025, GEICO was generating over **\$3 billion in annual underwriting profit** for Berkshire, making it one of the most valuable insurance companies in the world.

The key insight: **a cost advantage is the most durable moat when it is structural**. GEICO's advantage was not about being smarter or working harder. It was about having a distribution model that competitors could not copy without destroying their existing business. Traditional insurers could not simply fire their agents — their entire business was built around them.

Cost Advantage in AI

In AI, the most obvious cost advantage belongs to companies that can provide computing power cheaper than anyone else. This is why **Nvidia** has a moat that goes beyond just making good chips. Nvidia's CUDA software platform has trained **over 6 million developers** to use Nvidia chips exclusively. Writing AI models for AMD or Intel hardware would require retooling entire engineering teams. The cost of switching — in time, money, and risk — is so high that most companies simply do not do it.

Here is the data that proves the cost advantage:

Year	Nvidia Revenue	Gross Margin	CUDA Developers	AMD Gross Margin	Hyperscaler Combined Capex
2020	\$10.9 B	62%	~2 M	~44%	~\$120 B
2022	\$26.9 B	65%	~3 M	~45%	~\$200 B
2024	\$60.9 B	75%	~5 M	~51%	~\$251 B
2025	\$130.5 B	75%	~6 M	N/A	~\$416 B
Growth	12x in 5 years	13 pts (62% to 75%)	3x in 5 years	7 pts (44% to 51%)	3.5x in 5 years

The table reveals three things:

First, Nvidia's gross margin. At 75%, it is significantly higher than AMD's 51% or Intel's 35%. A gross margin this high means Nvidia can charge a massive premium for its chips — and customers still pay. Why? Because the CUDA software ecosystem locks them in. Once a company has trained its AI models on Nvidia hardware, switching to AMD means restarting from scratch. The cost of switching is higher than the premium Nvidia charges.

Second, the CUDA developer base. It has grown from 2 million to over 6 million in just five years. Every one of those developers is trained on Nvidia's platform. They write code for Nvidia. They optimize for Nvidia. They build careers around Nvidia. A competitor like AMD cannot simply offer a better chip — it would need to retrain millions of developers, rewrite thousands of AI models, and rebuild the data center infrastructure designed around Nvidia hardware.

Third, the hyperscaler capex. By 2025, the four largest cloud companies — Amazon, Microsoft, Google, and Meta — were collectively spending over **\$416 billion per year** on AI infrastructure. Amazon alone planned \$200 billion for 2026. A startup trying to compete on AI compute cost would need tens of billions of dollars and years of construction just to get to the starting line. The hyperscalers buy hardware in such enormous quantities that their per-unit costs are 30% to 50% lower than any

competitor.

Moat #2: Brand and Intangible Assets — The Coca-Cola Story

In 1988, Buffett began buying Coca-Cola stock. He eventually purchased \$1 billion worth — at the time, the largest single stock purchase Berkshire had ever made.

Critics said he was paying too much. The stock was trading at about 15 times earnings, which was expensive by 1980s standards. Buffett's response: *"Try to imagine a competitor raising \$100 billion and trying to take market share from Coca-Cola. It cannot be done. The brand is irreplaceable."*

He was right. Coca-Cola's brand is estimated to be worth over **\$100 billion** — more than the company's entire market value when Buffett bought it. The brand allows Coca-Cola to charge a premium price for sugar water. It allows the company to sell in every country on earth. And it allows the company to maintain its market share decade after decade, regardless of economic conditions.

Let us look at the numbers. When Buffett began buying Coca-Cola in 1988, the company had been through a rough patch — the infamous "New Coke" fiasco of 1985 had damaged management's credibility, and the stock market crash of 1987 had created a buying opportunity. Buffett invested approximately **\$1.02 billion** over 1988 and 1989, acquiring roughly 7% of the company. It was the single largest stock purchase Berkshire had ever made.

Here is what happened next:

Year	Revenue	Net Income	Stock Price
1988	~\$8.3 B	~\$1.0 B	~\$2.70
1993	\$14.0 B	\$2.2 B	~\$5.20
1998	\$18.8 B	\$5.0 B	~\$17.00
2005	\$23.1 B	\$4.9 B	~\$10.50
2010	\$35.1 B	\$11.8 B	~\$15.50
2015	\$44.3 B	\$7.4 B	~\$20.00
2025	\$47.9 B	\$13.1 B	~\$35.00

Buffett bought ~\$1.02B in 1988–1989

Growth	~5.8x	~13x	~13x
--------	-------	------	------

Year	Revenue	Net Income	Stock Price
------	---------	------------	-------------

The table tells a remarkable story. While Coca-Cola's revenue grew nearly six times over 37 years, its net income grew thirteen times — because the brand allowed it to raise prices faster than costs. Earnings per share grew eighteen times, boosted by aggressive share buybacks that the brand's consistent cash flows made possible.

By 2025, Berkshire's original \$1 billion investment was worth approximately **\$26 billion** in market value alone. Each year, those shares paid Berkshire roughly **\$776 million in dividends** — meaning the original investment effectively paid for itself every 1.3 years through dividend income alone. The total return, including dividends, was approximately **3,500%**.

Buffett has never sold a single share.

The See's Candy story is the same principle in miniature. Buffett later said: *"If I had bought See's at a price that reflected its full value, it would still have been a good investment because the moat was so wide. The brand was the moat."*

Brand Advantage in AI

Brand moats in AI are weaker than in consumer goods, but they exist. The strongest brand moat belongs to **Apple**. People do not buy iPhones because of the specifications. They buy them because they trust the brand, they love the ecosystem, and they are unwilling to leave. When Apple adds AI features to its devices, those features are not necessarily the best in the world — but they are trusted, seamless, and built into a product that a billion people already own.

Microsoft has a different kind of brand advantage. Its enterprise customers trust Microsoft with their most critical infrastructure. When Microsoft says its AI tools are enterprise-grade, secure, and reliable, CIOs believe it. A startup selling the same AI tool would face months of security reviews and trust-building. Microsoft starts every conversation with trust already established.

Moat #3: Switching Costs — The Microsoft Story

Switching costs are the most painful moat for competitors to overcome. The logic is simple: once a customer has invested time, money, and training into a product, the cost of leaving is too high — even if a competitor offers a better product.

Microsoft Office is the classic example. For thirty years, Microsoft Office has dominated the productivity software market. Not because it is the best product — there have been many technically superior alternatives. But because companies have spent **decades** training their employees on Word, Excel, and PowerPoint. Documents, spreadsheets, and presentations created in Microsoft formats are the standard

across industries. Switching would require retraining every employee, converting every file, and risking data corruption.

Here is what the switching-cost moat looks like in the numbers:

Metric	2015	2020	2023	2025	The Moat Story
Total Revenue	\$93.6 B	\$143.0 B	\$211.9 B	\$281.7 B	Nearly 3x in a decade — the moat compounds
Microsoft Cloud Revenue	~\$20 B	~\$60 B	~\$111 B	\$168.9 B	The toll road of the AI age
Azure Revenue	~\$8 B	~\$30 B	~\$56 B	\$75 B+	Fastest-growing part of the moat
M365 Commercial Seats	~150 M	~258 M	~382 M	450 M+	Every seat is a switching cost
M365 Consumer Subscribers	~20 M	~45 M	~80 M	89 M	Consumers locked in too
Stock Price	~\$46	~\$222	~\$376	~\$430	The market recognizes the moat

The table tells the story of a moat that has been **widening for a decade**. Every row grows every year — and the most important row is the installed base: over **450 million commercial seats** by 2025. Every one of those seats represents a company that has trained its employees on Microsoft software, built its documents and workflows around Microsoft formats, and integrated Microsoft into its operations. Leaving Microsoft is not a product decision — it is a corporate reorganization.

By 2023, Microsoft was generating over **\$45 billion per year** from its Office and cloud productivity suite — and that was before the AI boost. Now, with Microsoft Copilot integrated into every Office product, the switching costs have only increased. A company that leaves Microsoft Office does not just lose its spreadsheets — it loses its AI assistant, its embedded workflows, and its integration with the rest of its technology stack.

Switching Costs in AI

The strongest switching cost moat in AI belongs to **Nvidia**. When a company builds its AI models on Nvidia's CUDA platform, it is not just buying chips — it is adopting an entire software ecosystem. Developers learn to program in CUDA. Engineering teams optimize their models for Nvidia hardware. Data centers are designed around Nvidia's specifications.

A competitor like AMD cannot simply offer a better chip. It would need to retrain millions of developers, re-optimize thousands of AI models, and rebuild the data center infrastructure that supports Nvidia

hardware. The switching cost is so high that Nvidia can charge a premium — and customers still stay.

The same logic applies to cloud platforms. A company that builds its AI applications on **Microsoft Azure** using OpenAI's models cannot easily switch to Google Cloud. The engineering work is different. The APIs are different. The data pipelines are different. The switching cost is measured in millions of dollars and years of engineering time.

Moat #4: Network Effects — The American Express Story

Buffett has owned American Express stock since the 1960s, through multiple near-death experiences. The reason is simple: American Express has the strongest network effect in financial services.

A network effect exists when a product or service becomes more valuable as more people use it. Each new user adds value for every existing user. American Express is valuable to merchants because cardholders carry it. It is valuable to cardholders because merchants accept it. The more people who carry the card, the more merchants want to accept it — and the more merchants accept it, the more people want to carry it. This creates a self-reinforcing cycle that is almost impossible for a competitor to break.

By 2025, American Express had over **140 million cards in circulation** and processed over **\$1.5 trillion in transaction volume annually**. Competitors like Discover have tried for decades to replicate the Amex network effect and have never come close. The network creates the moat, and the moat protects the network.

Here is what that network effect looks like in the financial data:

Year	Card Members	Transactions Processed	Revenue	Net Income
1963 (Buffett buys)	~1 M	—	~\$0.1 B	~\$11 M
1997	~43 M	~\$250 B	~\$17.2 B	~\$1.4 B
2010	~88 M	~\$750 B	~\$30.0 B	~\$4.0 B
2025	~142 M	~\$1.7 T	~\$72.2 B	~\$10.8 B
Growth	142x	nearly 7x since 1997	~722x	~981x

The data reveals a beautiful compounding machine. As cards in force grew from 43 million to 142 million, transactions processed grew nearly seven times — because each additional cardholder also spent more over time. The network effect meant that every new cardholder made the service slightly more valuable for every existing cardholder (more merchants accepted Amex), which in turn made it easier to

sign up the next cardholder.

American Express had its challenges — the 2008 financial crisis nearly broke the company, and it had to convert to a bank holding company to survive. But the network effect was so strong that the business bounced back stronger than ever. By 2025, net income had reached a record **\$10.8 billion**.

Network Effects in AI

Network effects in AI are still emerging, but there are clear examples.

Microsoft + OpenAI is building a network effect. The more developers who build applications on OpenAI's models (via Azure), the more data OpenAI collects to improve its models, and the more developers want to use those improved models. This flywheel is real — but it is early. Whether it becomes a durable moat depends on whether the best AI models remain exclusive to a single platform.

Meta has a different kind of network effect. Its platforms — Facebook, Instagram, WhatsApp, Messenger — have over **3 billion daily active users**. Any AI feature Meta adds to these platforms instantly reaches billions of people. A startup cannot replicate this distribution advantage, no matter how good its AI model is.

Google has a search-based network effect. Every search query trains Google's AI. The more people who search, the better Google's search results become. The better the results, the more people who search. This data flywheel has been running for over twenty years and is one of the most durable competitive advantages in technology history.

Moat #5: Scale / Efficiency — The Burlington Northern Story

In 2010, Buffett bought Burlington Northern Santa Fe (BNSF) railway for \$34 billion — Berkshire's largest acquisition at the time. It was a bet on a business that had not changed in a hundred years.

Railroads are the ultimate scale business. Building a railroad costs billions of dollars per mile and requires government approval. Once a railroad is built, no competitor is going to build another one next to it. The cost of entry is so high that the existing players face almost no competition.

BNSF operates over **32,000 miles of track** across the western United States. It employs over 40,000 people. A competitor trying to replicate this network would need to spend over \$200 billion and navigate decades of regulatory approvals. It simply will not happen.

Here is what the scale moat looks like in the numbers:

Metric	2010 (Acquisition)	2025	The Moat Story
Track Miles	32,000	32,000	Cannot be replicated — no one will build a second railroad
Employees	~40,000	~40,000+	Operating scale that compounds
Revenue	~\$16.6 B	~\$23.4 B	Steady growth through every cycle
Net Income	~\$2.5 B	~\$5.5 B	More than doubled since acquisition
Dividends Paid to Berkshire	—	\$4.4 B	The moat pays cash every year
Purchase Price	\$34 B (13.6x earnings)	—	Now equal to only ~6x annual earnings

The table tells the story of a moat that is physically impossible to attack. The track miles do not change — because nobody can build a competing railroad. But the earnings grow every year, because the American economy grows, and everything the economy moves — coal, grain, containers, cars — moves on BNSF's fixed network. The \$34 billion purchase price looked expensive in 2010; by 2025 it represented barely six years of the company's earnings.

Buffett described BNSF as "*a bet on the future of the American economy*" — and it paid off. By 2025, BNSF was generating **\$5.5 billion in annual net income** for Berkshire.

Scale Advantage in AI

The best example of a scale moat in AI is **Amazon Web Services (AWS)**. AWS spends over **\$70 billion per year** on capital expenditures — building data centers, purchasing hardware, and laying fiber. No startup can match this. Even a well-funded competitor would need to spend a decade and hundreds of billions of dollars to build the global infrastructure that AWS already has.

Microsoft and Google are in the same race. By 2025, each of the three hyperscalers was spending over \$50 billion annually on AI infrastructure. The scale advantage means that a startup trying to offer cloud AI computing at competitive prices would need to spend money it does not have, on infrastructure it has not built, to win customers who already trust the incumbents.

This is the strongest moat in AI. It is not about technology — it is about **who has the checkbook large enough to stay in the game**.

The Buffett Moat Test: Six Questions

Before investing in any AI stock, Buffett would run the company through his moat test. Here it is, adapted for AI:

#	Question	What It Tests	Passes	Fails
1	Can competitors easily replicate what this company does?	Cost / Brand / Scale Moat	Nvidia (CUDA cannot be copied)	C3.ai (anyone can build enterprise AI apps)
2	How painful would it be for customers to leave?	Switching Cost Moat	Microsoft (CIOs locked into Azure ecosystem)	SoundHound (voice AI is replaceable)
3	Does the product get better the more people use it?	Network Effect Moat	Meta (billions of users)	Palantir (limited government contracts)
4	Can this company raise prices without losing customers?	Brand Moat	Apple (pricing power on every product)	AMD (chips are commoditized vs. Nvidia)
5	Is the competitive advantage getting stronger or weaker?	Moat Direction Test	Microsoft (AI makes Office stickier than ever)	IBM (cloud disrupted the legacy moat)
6	Would you bet the company's survival on this moat?	The Ultimate Test	Nvidia (yes, CUDA + hardware lead is definitive)	Most AI startups (no, unclear differentiation)

Fake Moats in AI: What Buffett Would Avoid

The most dangerous investments in AI are companies that **look** like they have moats but do not. Buffett has been fooled before (Dexter Shoes, IBM). Learn from his mistakes.

The "First Mover" Trap

Being first to market is not a moat. It is a head start. In AI, dozens of companies reached the market first with promising technology — and were quickly overtaken by better-funded, better-managed competitors.

AltaVista was the dominant search engine of the 1990s. It indexed the web faster and more completely than anyone else. In 1999, CMGI bought it from Compaq for **\$2.3 billion** — believing its first-mover advantage was unassailable. Four years later, AltaVista was sold to Overture for just **\$140 million** — a loss of **94%**. Google, which arrived later but built a real moat (PageRank algorithm, advertising network, data flywheel), had already made AltaVista irrelevant.

MySpace was the first social network to reach mass adoption. In 2005, News Corp bought it for **\$580 million** at the peak of its popularity. Six years later, it was sold for **\$35 million** — a loss of **\$545 million**, or **94%** of its purchase price. Facebook arrived later, but MySpace had no moat. Users could leave whenever they wanted — and they did.

Palm invented the smartphone category with the Palm Pilot and later the webOS platform. HP bought it in 2010 for **\$1.2 billion** believing the first-mover advantage in mobile was defensible. Within 18 months, HP wrote down **\$1.67 billion** on the investment — essentially admitting the entire purchase was worthless. Apple and Android arrived later but built real moats (app ecosystems, developer networks, brand loyalty) that Palm never had.

The pattern is consistent: first-mover advantage without a structural barrier to entry is not a moat. It is a head start that will be erased.

The "Data" Trap

Many AI startups claim that their data is a moat — that they have unique data sets that competitors cannot replicate. This is rarely true. Most valuable data is either publicly available, easily licensed, or can be generated by a competitor with enough resources.

The most expensive example of this trap was **IBM Watson Health**. IBM spent over **\$4 billion** acquiring health data companies and building an AI platform that it claimed would revolutionize medicine. The pitch was that Watson had unique access to medical data that no competitor could match, combined with AI that could diagnose cancer better than human doctors.

It was a disaster. Watson's recommendations were often inaccurate, hospitals stopped using the platform, and by 2022 IBM had sold the entire division for parts — receiving just over **\$1 billion** for assets it had spent four times that amount building and acquiring. That is a loss of roughly **\$3 billion**.

The problem was that healthcare data was not a real moat. Other companies had similar data. The technology did not work as advertised. And the competitive advantage IBM thought it had built — the data — was neither unique nor defensible.

The exceptions are companies with genuine **data network effects** — where every user interaction generates data that improves the product for every other user. Google Search is the best example. But most AI startups do not have this. Their data is only a moat until a larger competitor decides to compete.

The "Talent" Trap

A team of brilliant AI researchers is not a moat. People leave. In Silicon Valley, they leave frequently. The average tenure at an AI startup is less than two years. **A moat that depends on three PhDs who could quit tomorrow is not a durable moat — if the competitive advantage walks out the door with the team, it was never structural, and structural advantages are the only ones**

Buffett trusts.

Builder.ai raised over **\$450 million** from venture investors including some of the world's most sophisticated firms, reaching a valuation of **\$1.5 billion** in 2023. The company claimed its AI platform could build software apps automatically. Investors bet on the founding team's credentials. But investigations later revealed that much of the platform's output relied on human developers in low-cost countries, not AI. When the truth came out, the company collapsed into bankruptcy — destroying over a billion dollars in investor capital.

Humane, the AI hardware startup founded by former Apple employees, raised **\$241 million** and was valued at nearly **\$1 billion**. Its product — an AI pin worn on clothing — was supposed to replace the smartphone. When reviewers got their hands on it, the device was widely panned as unusable. The company burned through its funding and was eventually acquired for parts at a fraction of its peak valuation. The talent moat — "former Apple engineers" — turned out to be worth far less than investors assumed.

The rule is simple: **a real moat does not depend on any single person**. Nvidia's moat does not depend on Jensen Huang leaving tomorrow. Microsoft's moat does not depend on Satya Nadella. Coca-Cola's moat does not depend on any single employee. If a company's competitive advantage disappears when three people walk out the door, it never had a moat to begin with.

Moat Scoring for Major AI Companies

Let us apply the framework. Here is how the major AI stocks score across the five moat types — a tool you can use to evaluate any AI company yourself.

Company	Cost Advantage	Brand / Intangibles	Switching Costs	Network Effects	Scale / Efficiency	Overall Moat
Nvidia	Strong (CUDA cost lock-in)	Strong (recognized AI leader)	Very Strong (CUDA ecosystem)	Moderate (developer ecosystem growing)	Strong (fabless model, TSMC partnership)	Wide
Microsoft	Strong (hyperscaler economics)	Very Strong (enterprise trust)	Very Strong (Office / Azure lock-in)	Growing (OpenAI ecosystem)	Very Strong (global cloud infra)	Wide
Alphabet (Google)	Strong (hyperscaler)	Very Strong (search)	Moderate (search is)	Very Strong (search)	Very Strong (global cloud)	Wide

Company	Cost Advantage	Brand / Intangibles	Switching Costs	Network Effects	Scale / Efficiency	Overall Moat
	economics)	dominance)	low-switch)	flywheel)	infra)	
Meta	Moderate	Strong (social platforms)	Moderate (users can leave)	Very Strong (3B+ users)	Strong	Narrow to Wide
Apple	Moderate (premium pricing)	Very Strong (brand loyalty)	Strong (ecosystem lock-in)	Strong (App Store effect)	Strong (supply chain scale)	Wide
Amazon (AWS)	Very Strong (scale economics)	Strong (enterprise trust)	Strong (AWS + data lock-in)	Moderate (AWS marketplace)	Very Strong (global infra)	Wide
AMD	Moderate	Weak (commodity brand)	Weak (no CUDA equivalent)	Weak	Moderate	Narrow
Palantir	Weak	Moderate (gov contracts)	Strong (integration lock-in)	Weak	Moderate	Narrow
C3.ai	Weak	Weak	Weak	Weak	Weak	No Moat
SoundHound	Weak	Weak	Weak	Weak	Weak	No Moat

The Exercise: Find the Moat

Before you invest in any AI stock, write down the answers to these three questions:

1. **Which of the five moat types does this company have?** (Be specific. "A good product" is not a moat.)
2. **How wide is the moat?** (Can a well-funded competitor cross it in three years? One year? Never?)
3. **Is the moat growing or shrinking?** (This is the most important question. A shrinking moat will not protect you no matter how wide it was yesterday.)

Keep these answers. When the stock drops 30% and you feel panic, re-read what you wrote. If the moat is still intact, the lower price is an opportunity, not a threat. If the moat has eroded, the lower price was a warning.

Chapter Key Takeaways

1. **The Miracle of Omaha is real and measurable** — 19.9% annualized over 60 years turned \$10,000 into roughly \$550 million. It was built on a few simple ideas applied relentlessly.
 2. **The economic moat is the foundation of every great investment** — a durable competitive advantage protects a business the way a medieval moat protected a castle.
 3. **There are five types of moats** — cost advantage (GEICO), brand (Coca-Cola), switching costs (Apple), network effects (American Express), and scale (Burlington Northern). A business can have more than one.
 4. **The direction of the moat matters more than its existence** — a moat that is eroding (IBM) destroys value even if it was once wide; a moat that is deepening (Apple, Microsoft) compounds.
 5. **Fake moats are the most dangerous traps** — first-mover status (AltaVista -94%, MySpace - \$545M, Palm -\$1.67B write-down), claimed data advantages (IBM Watson -\$3B), and talent-dependent teams (Builder.ai -\$1B) are not real moats.
 6. **In AI, the strongest moats belong to the toll roads** — Nvidia's CUDA ecosystem (75% gross margin, 6M developers) and the hyperscaler cloud oligopoly (\$416B combined capex) are the deepest defenses in the AI value chain.
 7. **A real moat does not depend on any single person** — if a company's advantage disappears when three people walk out the door, it never had a moat to begin with.
-

References

GEICO Stock Price Timeline

1. *"The Impact of GEICO on Benjamin Graham and Warren Buffett"*, Index Fund Advisors, 2020 — GEICO stock fell from \$61 (1972) to \$2 (1976). ifa.com/articles/impact-geico-benjamin-graham-warren-buffett-luck-gecko
2. *"Warren Buffett's Geico Adventure Part 3: The Turnaround"*, Yahoo Finance, 2020 — Buffett bought GEICO at slightly above \$2 per share in 1976. finance.yahoo.com
3. *"GEICO Went from Brink to the Bank"*, The Washington Post, August 1995 — Stock fell from \$61 to \$7 in 1975, then to \$2 in 1976. [washingtonpost.com](https://www.washingtonpost.com)
4. *"GEICO History"*, GEICO corporate website — 1948 Graham investment, 1951 Buffett's first purchase, 1996 full acquisition. [geico.com/about/corporate/history](https://www.geico.com/about/corporate/history)
5. Berkshire Hathaway Annual Report, 1995 — Acquisition of remaining 49% of GEICO for \$2.3

billion.

Coca-Cola Investment

6. *"Coca-Cola Revenue History"*, WallStreetZen — Annual revenue data 1991–2025. wallstreetzen.com/stocks/us/nyse/ko/revenue
7. *"Coca-Cola's stock has had a total return of 3,534% since Warren Buffett first invested in 1988"*, CNBC, April 2025. cnbc.com
8. *"Warren Buffett's Investment in Coca-Cola: A Case Study"*, Picture Perfect Portfolios, 2023 — EPS growth from \$0.153 (1987) to \$0.82 (1997). pictureperfectportfolios.com

American Express

9. *"American Express — Financial History"*, Wikipedia — Annual revenue and net income table 2009–2025. en.wikipedia.org/wiki/American_Express
10. *"American Express Revenue History"*, WallStreetZen — Revenue data 1997–2025. wallstreetzen.com/stocks/us/nyse/axp/revenue
11. American Express 10-K SEC Filings, 1997–2025 — Cards in force and transaction volume data.

Microsoft Switching-Cost Moat

12. *"Microsoft Annual Report 2025"*, Microsoft Investor Relations — Total revenue \$281.7B, Microsoft Cloud \$168.9B, Azure \$75B+, M365 Commercial seats 450M+. microsoft.com/investor/reports/ar25
13. *"Microsoft Revenue History"*, Macrotrends — Revenue data 2012–2025. macrotrends.net/stocks/charts/MSFT/microsoft/revenue
14. *"Microsoft 365 Exceeds 450 Million Commercial Paid Seats"*, Microsoft Tech Community, FY26 Q2 — Commercial seat growth data. techcommunity.microsoft.com
15. *"Office 365 Reaches 400 Million Users"*, Office 365 IT Pros, January 2024 — 382M paid seats (April 2023), 2.5M seats/month growth. office365itpros.com
16. Microsoft FY25 Q4 Earnings — Microsoft 365 Consumer subscribers 89.0M (2025).

Nvidia & AI Cost Advantage

17. *"NVIDIA Revenue and Gross Margin History"*, Macrotrends — Revenue and gross margin data 2015–2025. macrotrends.net/stocks/charts/NVDA/nvidia/revenue
18. *"NVIDIA Announces Financial Results for Fiscal 2025"*, NVIDIA Newsroom, February 2025 — Data center revenue of \$115.2 billion. nvidianews.nvidia.com
19. *"CUDA, the De Facto Standard of HPC?"*, ScaleLang, March 2026 — 5 million CUDA developers (June 2024), 80% market share. scale-lang.com
20. *"Follow the CAPEX 2025 Retrospective"*, Platformonomics, February 2026 — Hyperscaler capex: Amazon \$200B, Google \$91.5B, Microsoft \$118B, Meta \$72B in 2025. platformonomics.com
21. *"Tech Megacaps to Spend More Than \$300 Billion in 2025 to Win in AI"*, CNBC, February 2025

— Combined Big Tech AI capex estimates. [cnbc.com](https://www.cnbc.com)

Fake Moats Examples

22. *"Overture Services to Buy AltaVista for \$140 Million"*, The New York Times, February 2003 — CMGI bought AltaVista for \$2.3B in 1999, sold for \$140M in 2003. [nytimes.com](https://www.nytimes.com)
23. *"News Corp. Sells MySpace for a Mere \$35 Million"*, ABC News, June 2011 — News Corp bought MySpace for \$580M in 2005, sold for \$35M in 2011. abcnews.com
24. *"HP to Acquire Palm for \$1.2 Billion"*, HP Investor News, 2010 — HP wrote down \$1.67B on Palm within 18 months. investor.hp.com
25. *"How IBM's Watson Went from the Future of Health Care to Sold Off for Parts"*, Slate, January 2022 — IBM spent \$4B+ on Watson Health, sold for ~\$1B. slate.com
26. *"Builder.ai: Anatomy of a Half-a-Billion-Dollar Deception"*, Segler Consulting, 2025 — Builder.ai raised \$450M+, valuation \$1.5B, collapsed into bankruptcy. seglerconsulting.com
27. *"A Billion Dollar AI Startup Just Collapsed Spectacularly"*, Yahoo Finance, 2025 — Humane raised \$241M, product failed, sold for parts. finance.yahoo.com

Berkshire Hathaway & Buffett Core Data

28. Berkshire Hathaway Annual Reports (1965–2025) — See's Candy, GEICO, BNSF, Coca-Cola, American Express investment data.
29. Buffett, W. — *Berkshire Hathaway Annual Shareholder Letters*, 1965–2025.
30. *"BNSF Railway Company 2010 Form 10-K"*, SEC — Revenue \$16.6B, net income \$2.5B at acquisition. [bnsf.com](https://www.bnsf.com)
31. *"BNSF reports net income gains in Q4, full-year 2025"*, Progressive Railroading — Revenue \$23.4B, net income \$5.5B in 2025; \$4.4B dividends to Berkshire. progressiverailroading.com

End of Chapter 1

Next: Chapter 2 — The Margin of Safety: Buying Dollar Bills for Fifty Cents

3. How to Invest in AI Stocks — Surviving and Profiting

"When there's a gold rush, don't dig for gold. Sell shovels." — Attributed to **Mark Twain**

Opening: Buying and Holding Isn't Enough — AI Moves Every Week

Let's say you bought AI stocks based on Chapter 3. **Now what?**

The Real Story of Samuel Brannan: California's First Millionaire

In 1848, gold was discovered at the American River in California. Hundreds of thousands of people abandoned their homes and traveled west, dreaming of striking it rich.

But the first person to become **a millionaire from the Gold Rush** never dug for a single ounce of gold. His name was **Samuel Brannan**. [1](#)

The moment Brannan heard about the gold discovery, he didn't grab a shovel and join the crowd. Instead, he quietly bought up **every pick, shovel, and gold pan in the region** — stocking them in his general store. Then he walked through the streets of San Francisco, holding up a vial of gold dust, shouting: *"Gold! Gold! Gold from the American River!"*

When the Gold Rush exploded into a frenzy, every prospector had no choice but to buy their equipment from Brannan's store — at vastly inflated prices. A gold pan he bought for 20 cents sold for **\$15**. The result? While most prospectors went home broke, Brannan pulled in roughly **\$5,000 a day** (~\$150,000 in today's money) and became one of the most powerful men in California.

The Story of Levi Strauss: From Canvas to the World's First Jeans

Another famous story from the same era: a German immigrant named **Levi Strauss** arrived in California with rolls of heavy canvas, intending to sell it for wagon covers and tents.

But when he arrived, he discovered the miners weren't complaining about their shelters — they were complaining about **their pants**. Prospecting was destroying regular trousers in days. Kneeling in streams, bending in mineshafts — regular fabric couldn't handle it.

Strauss pivoted. He started making durable work pants, reinforced the pockets with copper rivets, and created the world's first jeans. The brand **Levi's** is still going strong today. [2](#)



Translating These Stories to the AI Era

History keeps telling us the same lesson: in the most competitive gold rushes, the people who get rich sustainably are the ones providing **what everyone desperately needs**.

- If today's AI app startups are the **"prospectors"** gambling on whether their app hits or flops...
- **NVIDIA** is Samuel Brannan — the monopolist selling picks and shovels (GPUs) that every AI company needs, regardless of who wins.
- **Vertiv (VRT)** and **Equinix (EQIX)** are Levi Strauss — providing the durable infrastructure (power, cooling, data center real estate) that every AI system depends on, every second of every day.

This is why the **Infrastructure Layer** is the safest, most capital-preserving strategy in the AI revolution.

In this chapter, I'll teach you **how to monitor AI stocks intelligently** — without staring at charts all day, and without needing a computer science degree.

17

Part 1: The AI Investor Calendar — What to Watch and When

The Main Event: NVIDIA Earnings

In the AI investing world, nothing matters more than **NVIDIA Earnings** every 3 months.

Let me give you a real example. In February 2024, investors worldwide were on edge waiting for NVIDIA's quarterly results. Two days before the announcement, the stock was up 5%, then down 3%, swinging wildly.

The night of the announcement — NVIDIA reported Data Center revenue of **\$18.4B** — up **409% YoY** — crushing analyst estimates.

The next morning, NVIDIA **surged 16% in a single day** — adding roughly **\$277 billion** in market cap in 24 hours (a US record at the time). AI stocks across the board — AMD, TSMC, Microsoft — all rallied with it. [3](#)

This is the power of NVIDIA Earnings on the entire AI ecosystem. It's the **Super Bowl of AI**.

Timing	Event	What to Watch
Mid Feb	NVIDIA Q4 Earnings	Data Center revenue, Gross Margin, Guidance
Mid Apr	NVIDIA Q1 Earnings	YoY growth — accelerating or decelerating?
Mid Aug	NVIDIA Q2 Earnings	Data Center revenue vs. estimates
Mid Nov	NVIDIA Q3 Earnings	Outlook for next year + Vera Rubin updates
Every Quarter	Microsoft, Google, Amazon Earnings	AI Revenue Run Rate, CapEx spending
Every Quarter	Vertiv, Equinix, Broadcom Earnings	Data center buildout trajectory

Other Key Events

Month	Event	Importance
March	NVIDIA GTC Conference	☆☆☆ — Jensen keynote, new product launches
June	Apple WWDC	☆☆ — AI on iPhone, Apple Intelligence updates
October	AWS re:Invent	☆☆ — AI cloud, new services
Quarterly	Fed Interest Rate Decisions	☆☆☆ — impacts all tech valuations

💡 **Put NVIDIA earnings dates on your calendar right now.** Nothing moves AI stocks more.

Part 2: 5 KPIs Every AI Investor Needs to Know

When you read AI news, you'll encounter a ton of technical jargon. But as an investor, **you don't need to know everything**. Just these 5 KPIs:

KPI #1: AI Revenue Run Rate — Real Money or Just Hype?

In 2023, an AI startup called **Stability AI** had a \$1B valuation. Employees told investors "we are the future of AI." But when the numbers came out — their total 2023 revenue was roughly **\$8–11M** — against expenses of ~\$153M. [4](#)

This is textbook "fake it till you make it" — sky-high valuation, tiny real revenue. Sound familiar? It's the same story as **Pets.com** in the dot-com era: \$300M raised, cute TV commercials, bankrupt in 9 months.

The key metric: AI Revenue **must grow faster** than the core business:

Company	AI Revenue / Run Rate	What It Means
Microsoft	>\$37B/year	~11% of total revenue — real and growing
Google Cloud	\$20B/quarter (+63%)	AI is accelerating cloud growth
Amazon AWS	AI services growing fast	Enterprise customers adopting AI tools

🔑 If AI revenue is growing slower than the core business — the adoption narrative is broken.

KPI #2: CapEx — Are Big Tech Companies Daring to Spend?

In 2024, Mark Zuckerberg announced massive CapEx plans for AI data centers (roughly **\$35–40B**). Investors panicked — Meta dropped 10–15% in a day, afraid Meta was overspending.

But Meta's AI (Llama) went on to become one of the most-used open-source models in the world, and Meta hit new all-time highs. [5](#)

Company	CapEx Trend	Signal
Microsoft, Google, Amazon	Increasing	● They believe AI will grow
NVIDIA	R&D increasing — Vera Rubin	● Maintaining leadership
Companies cutting CapEx	Decreasing	● May be losing confidence

KPI #3: NVIDIA Gross Margin — The Mother of All AI Metrics

NVIDIA's gross margin is one of the most important signals in the entire AI ecosystem:

Fiscal Year	Gross Margin	The Story
FY2022	64.9%	Baseline before AI boom
FY2023	56.9%	⚠️ Panic! — Inventory write-down (gaming slump + crypto crash)
FY2024	72.7%	● AI chip demand surged — NVIDIA gained pricing power
FY2025	~75%	👑 Highest in company history
FY2026	71.1%	Vera Rubin ramping — margins slightly normalized

[6](#)

🔑 **Margin > 75%** = NVIDIA is dominant, competitors can't catch up **Margin < 65%** = Competition is pressuring, or customers are negotiating harder

KPI #4: AI Adoption Rate — Are People Actually Using AI?

ChatGPT went from 0 to 1 billion users in 3.5 years — faster than any app in history:

Metric	Latest Figure
ChatGPT Weekly Active	900 million

Metric	Latest Figure
--------	---------------

ChatGPT Monthly Active ~1 billion

Enterprise AI Adoption ~72% of companies using AI

🔑 If adoption slows for 2 consecutive quarters — it may be time to reduce exposure.

KPI #5: AI Patent Count

More AI patents = more investment in the space. Fewer = potential saturation.

🚩 Part 3: Warning Signs — When to Sell

Knowing **when to sell** is just as important as knowing when to buy.

Let me tell you about the **SoftBank Vision Fund**. In 2017, SoftBank invested over **\$100B** in tech companies, including WeWork (\$4.7B), Uber, and DoorDash. By 2019, warning signs were everywhere — WeWork's IPO collapsed, its valuation dropped 80%. But Masayoshi Son refused to sell. The longer he held, the more he lost.

By 2020, SoftBank reported a **\$17.7B loss** from the Vision Fund alone — one of the largest losses in Japanese corporate history. [Z](#)

Lesson: When you see warning signs — act.

🔴 Level 1 — Watch Closely (Don't Sell Yet, But Be Cautious)

Signal	Real Example	What It Means
NVIDIA guidance miss 2 quarters in a row	Similar to 2022 — NVIDIA missed guidance 3 quarters straight — stock dropped 66%	AI demand may be slowing
Big Tech cuts data center CapEx	If Microsoft stops building data centers — AI hype could be ending	Confidence declining
AI startups failing en masse	2023 saw ~20% of AI startups fail — still not peak	Bubble beginning to crack
Strict AI regulation	EU AI Act could limit certain capabilities	Growth could slow

🔴🔴 **Level 2 — Should Reduce Positions**

Signal	Action
NVIDIA Margin < 65%	Reduce infrastructure weight
Fed raises rates aggressively	Tech stocks correct — AI isn't immune
Entire sector overvalued	Take profits
Adoption slows 2 quarters in a row	Step back, wait

🟢 **Green Signals — When to Add**

Signal	Action
AI agents adopted by enterprise	Increase Application Layer exposure
Vera Rubin exceeds expectations	Buy more NVIDIA
Major corporation announces AI data center buildout	Add VRT, EQIX
20%+ correction on solid company	DCA — buy when others are fearful

🔧 **Part 4: Tools + 30 Minutes Per Week**

US Tool	What It's For	Price
Yahoo Finance	Stock prices, news, financials for NVDA, MSFT, etc.	Free
SEC.gov (EDGAR)	Read 10-K, 10-Q filings — the real numbers	Free
Finviz	Stock screener, heat maps, sector performance	Free / \$39.50
TradingView	Charts, technical analysis, community ideas	Free / paid tiers
Bloomberg / CNBC	Breaking news, earnings coverage	Free / cable

My Weekly 30-Minute Routine:

📅 AI WEEKLY CHECKLIST 

| 🕒 5 min: Check your portfolio prices |

	🕒 10 min: Read this week's AI news	
	🕒 5 min: Check 1 KPI	
	🕒 5 min: Review warning signs	
	🕒 5 min: Mindset reset (especially	
	if you're panicking)	
	Total: 30 min/week	

📖 Part 5: Mindset — Lessons from People Who Got Rich (and People Who Got Wrecked)

The most important thing isn't tools or data — it's **mindset**.

1. 🧘 AI = 10 Years, Not 10 Days

The Real Winner of the Dot-Com Era: Cisco Systems

During the dot-com boom of 1995–2000, hundreds of new companies jumped into the gold rush.

Pets.com raised \$300M, had cute TV ads, and went bankrupt 9 months after its IPO.

But amid the thousands of companies that died, one quietly played the **pick-and-shovel** role and got rich without competing with anyone: **Cisco Systems**.

Cisco didn't build websites. It built **routers and switches** — the hardware that made the Internet work. Their logic was simple but powerful:

"No matter which dot-com company wins... every single one needs to build a website, and to make that website work, they need to buy Cisco networking equipment first."

While prospecting companies were going bankrupt, Cisco's revenue surged from \$1.3 billion in 1994 to **\$18.9 billion in 2000**. It briefly became the most valuable company in the world. [9](#)

🤖 Applying This to Your AI Portfolio

- **Don't bet on AI apps yet:** Today's flashy AI apps (image generators, writing bots) are Pets.com — exciting, easy to copy, and ready to die.
- **Invest in infrastructure monopolies:** **NVIDIA** (GPUs), **Vertiv** (cooling/power), **Equinix** (data centers) — these are **Cisco in the making**.

No matter which AI app wins, every single one needs GPUs, data centers, power, and cooling. The infrastructure layer gets paid either way.

2. 📉 Corrections Are Normal — Like Breathing In and Out

Here's the history of NVIDIA corrections (approximate peak-to-trough): [8](#)

Year	What Happened	Drop	Came Back?
2018	Crypto crash — Jensen said "crypto contributed nothing" — investors panicked	-54%	✅ ATH 2020
2020	COVID — global markets shut down — "next Great Depression"	-32%	✅ ATH 2021
2022	Fed rate hikes — NVDA from \$330 to \$100 = -66% — "AI is a bubble"	-66%	✅ ATH 2024
2024	Competition fear — AMD launched MI300 — "NVIDIA has a real rival"	-35%	✅ ATH 2025
2025	Tariff war — supply chain fears	-45%	✅ ATH 2026

Every drop had scary news. If you waited for good news to buy, you'd buy at ATH (All Time High) every single time.

If you believe in the trend — buy the dips or dollar-cost average.

3. 🎯 Never Go All-In on One Theme

Flashback to 1999. The dot-com mania was at its peak. Retail investors everywhere were going **all-in** on Internet stocks — Pets.com, Webvan — seduced by the hype and the "this changes everything" narrative, ignoring that these companies had zero profits.

Meanwhile, smart institutional money did the opposite. They pulled 60–70% of their portfolio out of speculative app-layer stocks and parked it in **infrastructure** — specifically **Oracle** (databases that every website needed) and **Microsoft** (the operating system every computer ran on).

Their logic: no matter which online retailer wins or dies, they all *need* Oracle databases and Windows.

When the dot-com crash came (Nasdaq -78%), the app-layer prospectors went to zero. But Oracle and Microsoft — they survived, recovered, and went on to dominate the next two decades.

Lesson: No matter how much you believe in a theme — **never go all-in.**

4. 📖 Keep Learning

AI changes fast. But you don't need to understand the code — just the business impact.

5. 🧠 Consult Professionals When in Doubt

If you're unsure — speak to a licensed financial advisor or start with an ETF.

🇮🇹 Will There Be Another AI Winter?

AI Winter is a period when AI hype dies down. It's happened **twice** before:

AI Winter #1 (1974–1980): The UK government commissioned Professor Lighthill to evaluate AI. His report concluded AI "couldn't solve real problems" — the government cut all AI research funding. AI researchers lost their jobs. The field stagnated for 6 years.

AI Winter #2 (1987–1993): A bubble formed around "Expert Systems" — hundreds of companies sold AI that "mimicked human experts." But the technology couldn't deliver on its promises. The bubble burst, companies failed, AI funding vanished almost entirely. [10](#)

But this time is fundamentally different:

Past AI Winters	This Time Is Different
AI had no real revenue	AI is generating >\$100B in revenue
Confined to universities	Deployed in real businesses — Microsoft, Google
No infrastructure	Billions in data centers already built
Zero user adoption	ChatGPT: 1 billion real users
Just research projects	Real products shipping — Copilot, Gemini

💡 **AI Winter is very unlikely this time** — AI is already generating massive real-world value, unlike past cycles where it was just a lab curiosity.

🎯 Summary: The AI Investor's Weekly Checklist

Every Sunday, spend just **30 minutes** on this:

1. **5 min:** Check portfolio prices
2. **10 min:** Read the week's AI news
3. **5 min:** Check 1 KPI (rotate through the 5)
4. **5 min:** Review warning signals
5. **5 min:** Mindset reset (especially if you're panicking)

⚠ **Disclaimer:** All content, data, and strategies in this chapter are for educational purposes only. They do not constitute investment advice. Buy/sell signals and KPIs are monitoring frameworks, not guarantees of results. Always do your own research and consult a licensed financial advisor before making investment decisions.

Footnotes & References

All data in this chapter verified as of July 2026 from the following sources:

Where these came from

The full books, all on Amazon:

Book	What it answers
A Few Stocks Move the World	Why a handful of companies now move entire markets — and what that means for an index fund
Buffett & Munger's Investment Wisdom	Applying value investing — moats, margin of safety, circle of competence — to the AI age
AI Stocks: The Stocks That Change Your Life	A data-driven playbook for the AI super-cycle: semiconductors, platforms, infrastructure
DRAGON CHECKMATE: The Race of AI Supremacy	The US–China AI contest, and how companies survive it
Bitcoin Alpha	How to size, custody and hold bitcoin as part of a portfolio
Money on the Mind	Why wealth is a mindset before it is a number

If a chapter here was useful, the book it came from goes further. And if you read one and it lands, an honest review helps the next reader find it.

© 2026 Victor Wong. All rights reserved. These chapters are taken from previously published works and are reproduced here for sampler purposes.

Questions, corrections or a book you think I should write: doctorwe@gmail.com

-
1. Samuel Brannan was "California's first millionaire." He bought gold pans for ~20 cents and sold them for \$15. His store grossed ~\$150,000/month (~\$5,000/day ≈ ~\$150,000/day in today's money). — Wikipedia; PBS American Experience; Legends of America; HowStuffWorks. The "sell shovels" quote is commonly attributed to Mark Twain but lacks verified sourcing. [↪](#)
 2. The Levi Strauss story (canvas tents → jeans) is popular lore. Historically, Strauss was a dry goods wholesaler who arrived in San Francisco in 1853. Riveted jeans were officially patented in 1873 with tailor Jacob Davis — Levi Strauss & Co. company history. [↪](#)
 3. NVIDIA Q4 FY2024 (reported Feb 21, 2024): Data Center revenue \$18.4B, +409% YoY. Stock surged ~16% the next day, adding ~\$277B in market cap (a US single-day record at the time) — NVIDIA SEC Form 8-K; CNBC. [↪](#)
 4. Stability AI raised \$101M at \$1B valuation (2022). 2023 revenue estimated at \$8–11M vs. ~\$153M in total expenses (revenue of ~\$1.5M was for 2022) — Forbes; Sacra; The Register. [↪](#)
 5. Meta's AI CapEx spending and market reactions — Meta was sold off multiple times on AI spending fears (e.g., Apr 2024 and Apr 2026 with ~10% drops) before recovering. Meta 2026 CapEx: ~\$125–145B — CNBC; Fortune; Bloomberg. [↪](#)
 6. NVIDIA GAAP Gross Margin by fiscal year: FY2022 64.9%, FY2023 56.9% (inventory write-down), FY2024 72.7%, FY2025 ~75.0%, FY2026 71.1% — NVIDIA SEC Form 8-K. [↪](#)
 7. SoftBank Vision Fund launched 2017 with ~\$100B. Reported ~\$17.7B loss for fiscal year ending Mar 2020 — SoftBank Group earnings; financial news reports. [↪](#)
 8. NVIDIA correction history approximate peak-to-trough — actual figures vary by measurement period (e.g., the 2025 cycle ranged from -37% to -45% depending on reference points) — market data. [↪](#)
 9. Cisco Systems revenue grew from ~\$1.24B (FY1994) to ~\$18.9B (FY2000) and became the most valuable company in the world in March 2000 — Cisco annual reports; market data. [↪](#)
 10. AI Winter #1 (~1974–1980): triggered by the Lighthill Report (1973) causing UK to cut AI research funding. AI Winter #2 (~1987–1993): collapse of the Expert Systems market and LISP machine industry — AI history sources (e.g., Crevier, "AI: The Tumultuous History"). [↪](#)
 11. ChatGPT ~900M weekly users (Feb 2026) — OpenAI. Enterprise AI adoption ~72% — McKinsey "The State of AI" (2024) survey. [↪](#)