Equities Outlook 2026

Investing in an evolution as disruptive as the automation of computing, while navigating an impression of the K-shaped economic backdrop

2026 Theme of Positioning:

Complexity within Simplicity

The economic environment through 2026 favors equity posture rooted in the foundations of the AI boom, duly positioned to capture value from both continued growth of key drivers and the alleviation of structural chokepoints within the broader boom. Under the easing macroeconomic conditions through 2026, respect to risk of technology-heavy concentration can be managed by allocating toward earnings-growth positioned names that offer insulation from nonessential civilian consumer-facing cyclicality. In this setting, alpha-enhanced portfolio construction provides a pathway to improve the efficiency of passive index allocations amid a complex and rapidly evolving K shaped impression on holistic market growth attributable to the participation of different subdivisions of the market itself.

Key Takeaways

- **Part 1: Bounds**
- **Economic Impressions**
  - K-Shaped Backdrop and Subsequent Impressions
  - The Structural Risk of the Civilian Consumer
  - Continuation
  - Alliance Dependent Fragmentation

- **Define a “Bubble”**
  - 1945 -1959 not 2000 - 2001
  - Behavioral Biases

- **Part 2: Artificial Intelligence**
- **An Evolution as Significant as the Rise of Computing Itself**
  - Cycles
  - Manufacturing Determines World Order, Technological Advancement Determines Manufacturing Dominance

- **The Dual Dependency Constraint Loop and The External Third Variable.**
  - Self-reinforcing Constraints
  - Thomas Kuhn’s Paradigm Shift Theory
  - The Third Variable

Part 1. Bounds

Economic Impressions

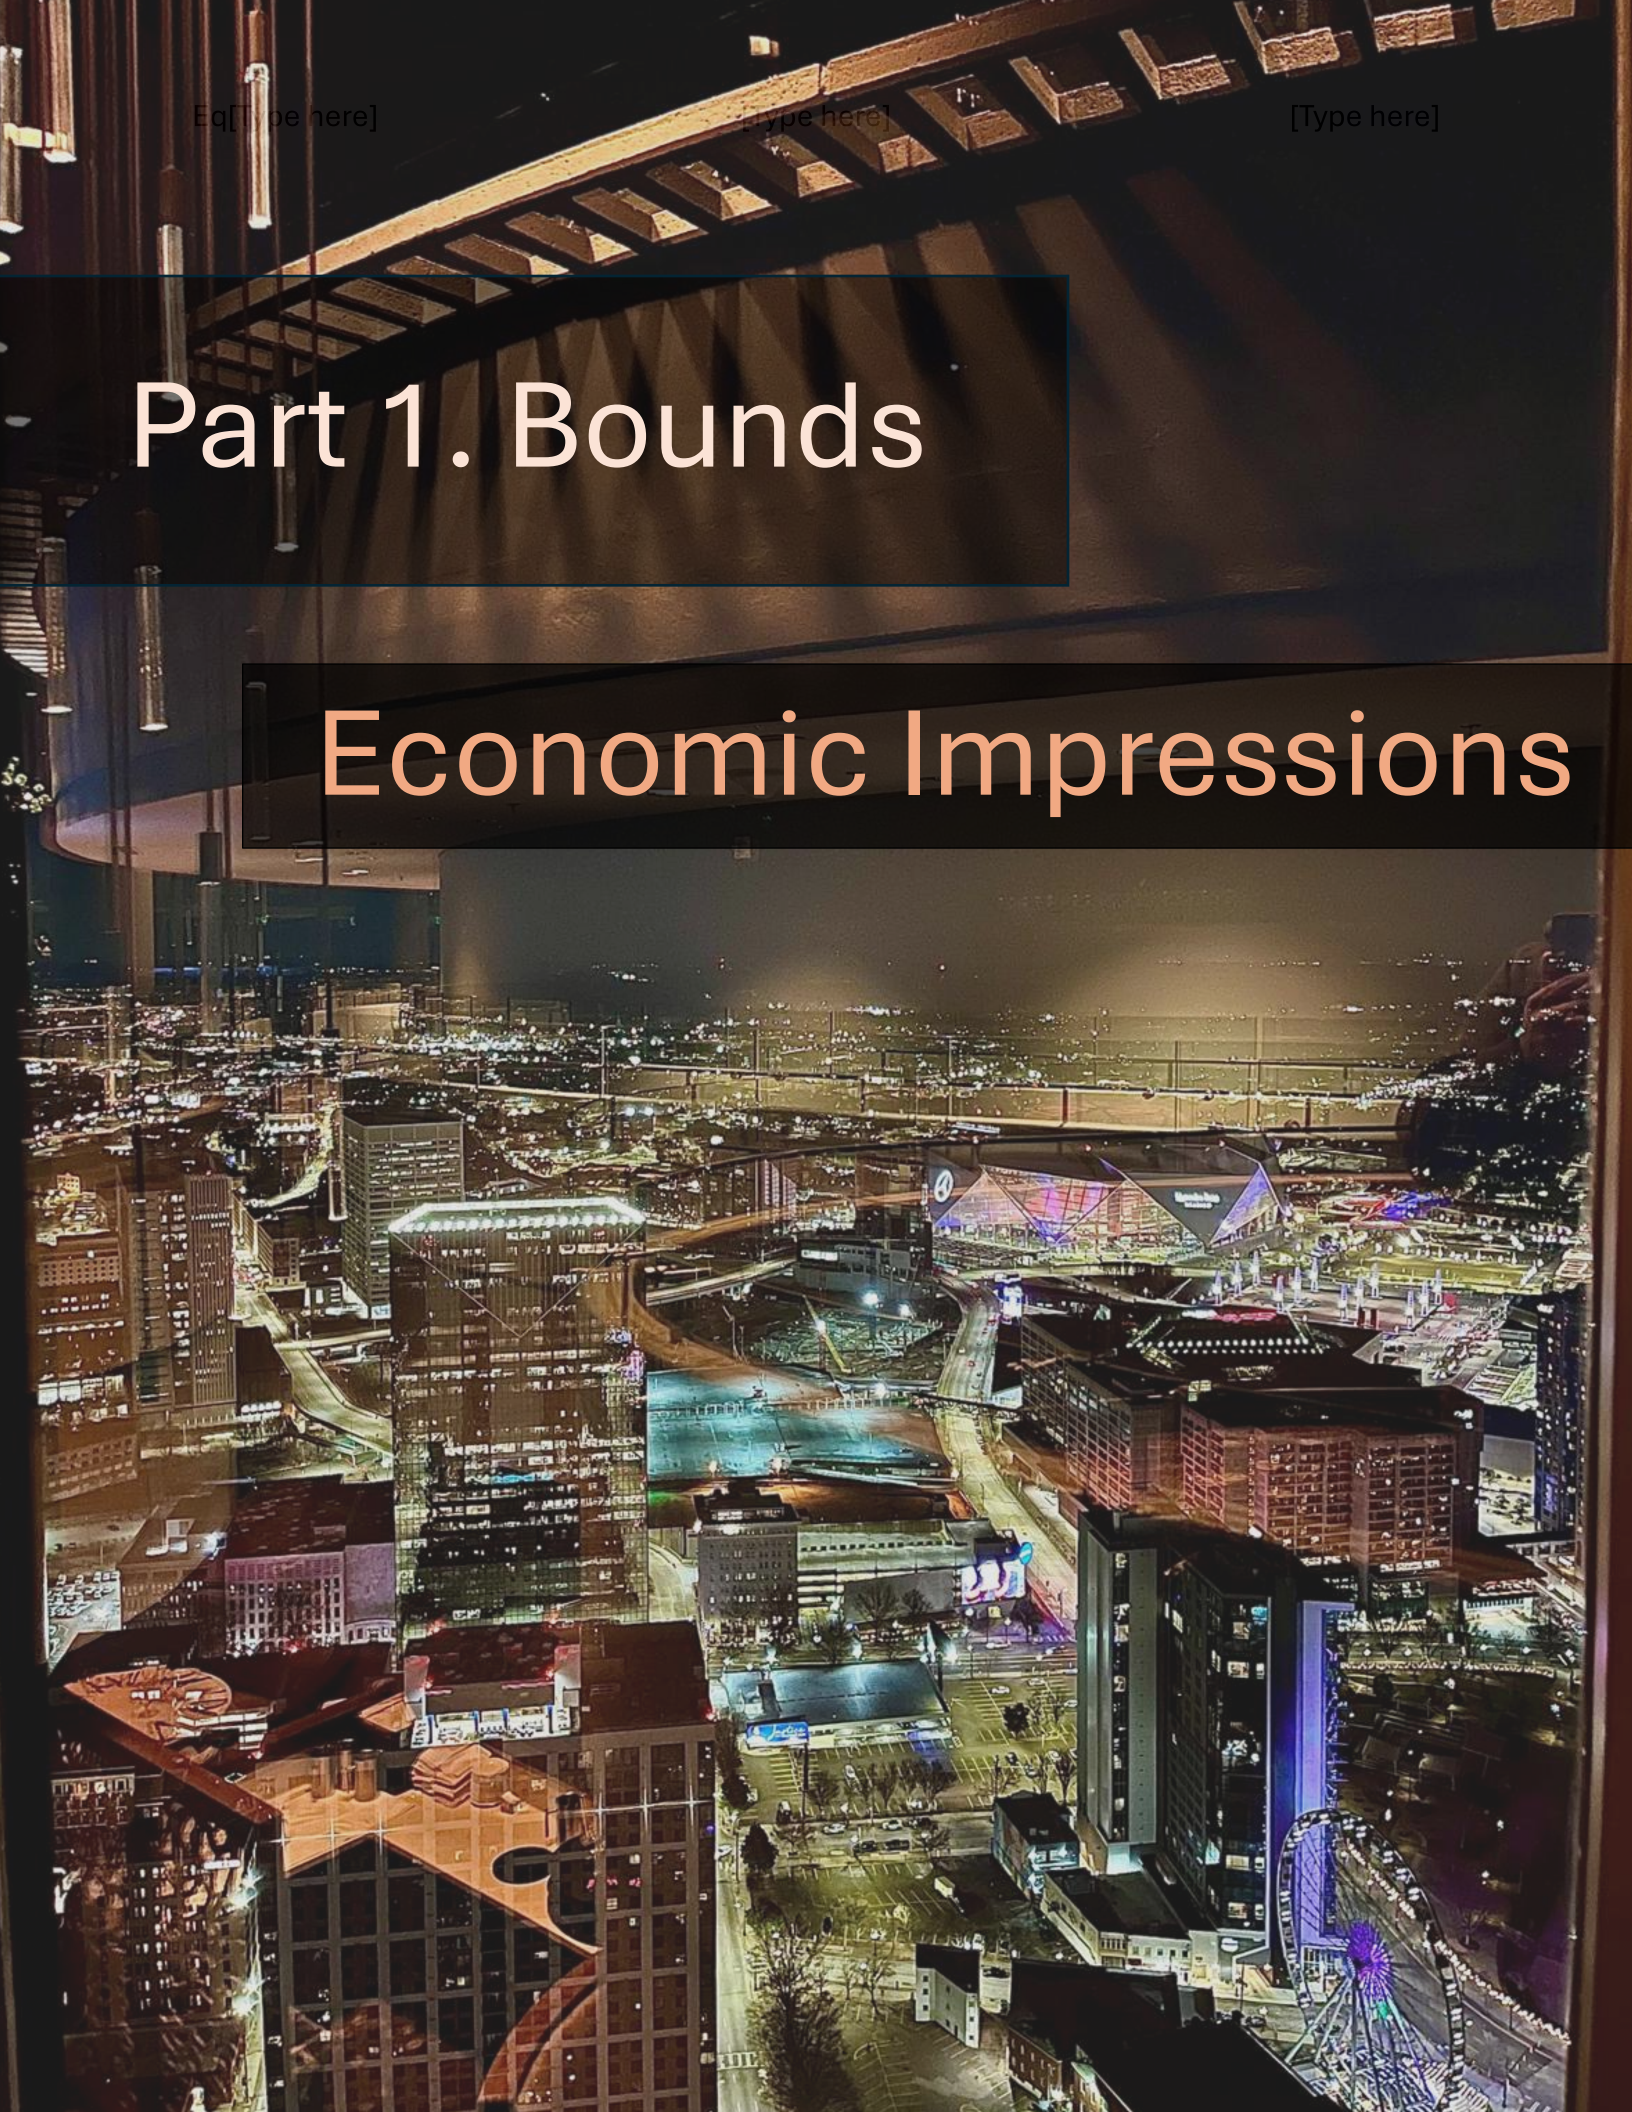

K-Shaped Backdrop and Subsequent Impressions

Underway is an era defined by higher structural inflation and an entrenched K-shaped economic profile, shaped by divergence across different cohorts within the American consumer base. The disparity across market participation that has persisted since the 2020 pandemic should intensify as broadening of participation formulates as an ideal corresponding only across individual names of existing dominant sectors while sector participation itself remains unequal as an explicit reflection of the K-shaped consumer base.

The current cycle positioning of Artificial Intelligence, and the supporting physical hardware therein, is consistent with earlier disruptive technological eras in which name-level broadening was common even as overall sector participation across the market stayed narrow.

Conversely, the lower arm of the K tells a different story. While policy dictated levels of the macroeconomic environment should improve, rather than a nod to progress, these improvements are a consideration of an era of higher structural inflation and the inability of current models to accurately align policy(s) with more volatile data.

Heading into 2026, the narrative facing the civilian consumer, and the business that faces it, is weak. Pyrite polished by the notion of rate cuts to be perceived as gold.

The Structural Risk of the Civilian Consumer

The current magnitude of the base effect of wealth, combined with the likelihood that translation from policy level rate cuts into meaningful reductions in consumer borrowing costs will remain delayed under an increasingly riskier consumer profile in addition to structurally higher and more volatile inflation, creates an environment in which the civilian consumer is unlikely to sustain the resilience observed since the 2020 pandemic thus far.

This view is reinforced by the reality that the consumer appears increasingly stretched, both through rising dependence on debt and through a post-Covid shift in consumer psychology that has favored short-term spending over long-term balance-sheet health.

Within this framework, the narrative that broadening in market participation will be driven by small and mid-caps becomes more complicated. The fundamental performance of these businesses appears constrained by the rate path alone, while performance tied to core business function, top-line growth, and true demand elasticity hinges critically on whether the stretched consumer can bend without breaking. Monetary relief may offer valuation support, but it does not necessarily repair the underlying demand profile.

Layered onto this challenge is an effective tariff rate above 18%, persistent fiscal deficits, and a household wealth distribution increasingly concentrated in illiquid or "un-tappable" forms. These factors reinforce the base effect dynamics that feed into the K-shaped divergence, further widening the gap between the lower arm of the K and the consumer cohorts that can participate in growth of the market as well as directly contribute to the business that heavily depends on civilian sales and engagement.

Continuation

Portfolios that have demonstrated high statistical efficiency over the past three years should continue pulling their weight through 2026–2027, without significant mean reversion. Under this lens of continuation, several foreign developed and emerging markets that have benefited from their role as critical chokepoints in U.S. technological infrastructure and materials supply chains are positioned to outperform returns of domestic indices in nearly every scenario in which domestic equity leadership remains disproportionately concentrated.

Spending from the primary leaders of the AI boom, companies such as Alphabet, Amazon, Meta, Microsoft, Oracle, and Nvidia, now account for nearly 25% of total U.S. capital expenditures.[1] This cohort contributed more to U.S. GDP growth in 2025 than consumer spending did. Yet even with this magnitude of investment, total AI-related capex still measures only around 1% of GDP while investment during previous disruptive technological cycles such as electricity, railroads, and telecommunications peaked between 2% and 5% of GDP.[2]

1. Empirical Research Partners. The Hyperscalers: *Making the Jump to Hyperspace?* August 11, 2025.
2. J.P. Morgan Asset Management – *Promise and Pressure*. Data as of December 17, 2025

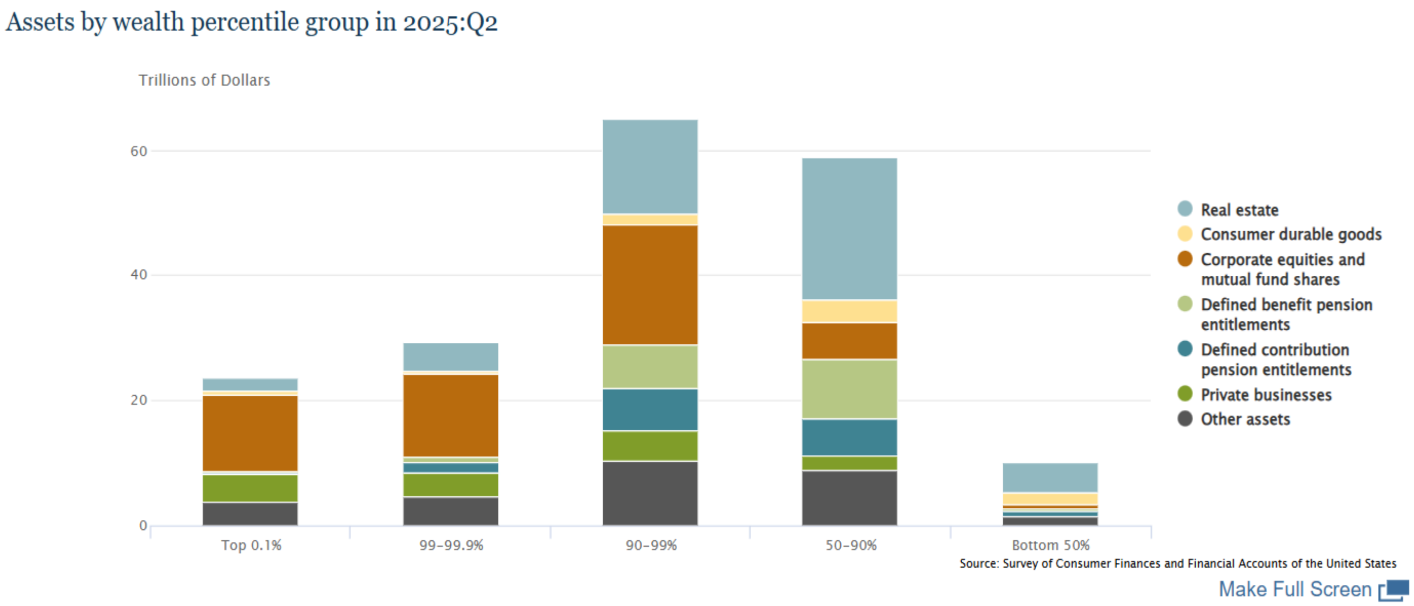

Wealth component	Top 0.1% (US$ Trillions)	99-99.9% (US$ Trillions)	90-99% (US$ Trillions)	50-90% (US$ Trillions)	Bottom 50% (US$ Trillions)
Real estate	1.97	4.66	15.10	22.73	4.85
Consumer durable goods	0.68	0.40	1.80	3.59	2.00
Corporate equities and mutual fund shares	12.27	13.28	19.11	5.99	0.54
Defined benefit pension entitlements	0.30	0.91	6.87	9.44	0.48
Defined contribution pension entitlements	0.22	1.56	6.78	6.06	0.72
Private businesses	4.32	3.88	4.90	2.26	0.17
Other assets	3.81	4.61	10.34	8.84	1.39

Board of Governors of the Federal Reserve System, Distributional Financial Accounts (DFA) and Survey of Consumer Finances (SCF), 2025 Q2.

The distribution of assets across wealth percentiles highlights the current extent of the base effect between the top decile of U.S. households and the remaining 90%. As shown, the top 10% holds a disproportionately large share of total household assets, with a significant concentration in corporate equities, mutual funds, and private business ownership. These asset classes are directly linked to financial market performance and tend to reprice continuously with changes in valuations.

In contrast, the bottom 90% holds a greater share of wealth in real estate, consumer durable goods, and pension entitlements. These assets are generally more illiquid, less responsive to short-term market movements, and less directly accessible for discretionary spending or balance-sheet adjustment. The data illustrates that exposure to financial assets declines sharply outside the top decile, while reliance on non-financial or contractual forms of wealth increases. This structural composition difference helps explain why changes in financial market conditions translate unevenly across household cohorts.

TOP 10% PERCENTILE CHANGE IN WEALTH ATTRIBUTABLE TO BOTH CHANGE IN ASSET OWNERSHIP % AND CHANGE IN PRICE OF S&P500 Y/Y

SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.945346352
R Square	0.893679725
Adjusted R Square	0.892127604
Standard Error	0.002724404
Observations	140

ANOVA

	df	*SS*	*MS*	*F*	*Significance F*
Regression	2	0.008547306	0.004273653	575.7797491	2.10476E-67
Residual	137	0.001016865	7.42237E-06		
Total	139	0.009564171			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*	*Lower 95.0%*	*Upper 95.0%*
Intercept	0.000978918	0.000271971	3.599350772	0.000444762	0.000441115	0.001516721	0.000441115	0.001516721
X Variable 1	-0.015567696	0.002065513	-7.536962512	5.92613E-12	-0.019652107	-0.011483286	-0.019652107	-0.011483286
X Variable 2	1.089622977	0.039039125	27.91105038	2.1982E-58	1.012425793	1.16682016	1.012425793	1.16682016

BOTTOM 50% PERCENTILE CHANGE IN WEALTH ATTRIBUTABLE TO BOTH CHANGE IN ASSET OWNERSHIP % AND CHANGE IN PRICE OF S&P500 Y/

SUMMARY OUTPUT

Regression Statistics	
Multiple R	0.434346182
R Square	0.188656606
Adjusted R Square	0.176812176
Standard Error	0.002300565
Observations	140

ANOVA

	df	*SS*	*MS*	*F*	*Significance F*
Regression	2	0.0001686	8.42999E-05	15.92787661	6.03285E-07
Residual	137	0.000725086	5.2926E-06		
Total	139	0.000893686			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*	*Lower 95.0%*	*Upper 95.0%*
Intercept	-0.000840657	0.000225682	-3.724954075	0.00028462	-0.001286928	-0.000394385	-0.001286928	-0.000394385
X Variable 1	0.007659332	0.001378388	5.55673033	1.38347E-07	0.004933664	0.010385	0.004933664	0.010385
X Variable 2	0.312759752	0.083888284	3.728288833	0.000281225	0.146876441	0.478643062	0.146876441	0.478643062

Jordan Wylie, *Wealth Gap Analysis 2025*. Wylie ThinkTank

An analysis of quarterly deviations in total wealth across cohorts compared to corresponding changes in liquid asset ownership as a share of total wealth and movements in the S&P 500 shows a pronounced divergence. For the top 10% of U.S. households, 89.21% of the variance in total wealth is explained by changes in collective liquid-asset ownership and the price of the S&P 500. In contrast, for the bottom 50% of households, only 17.68% of the variance in total wealth is attributable to these same factors.

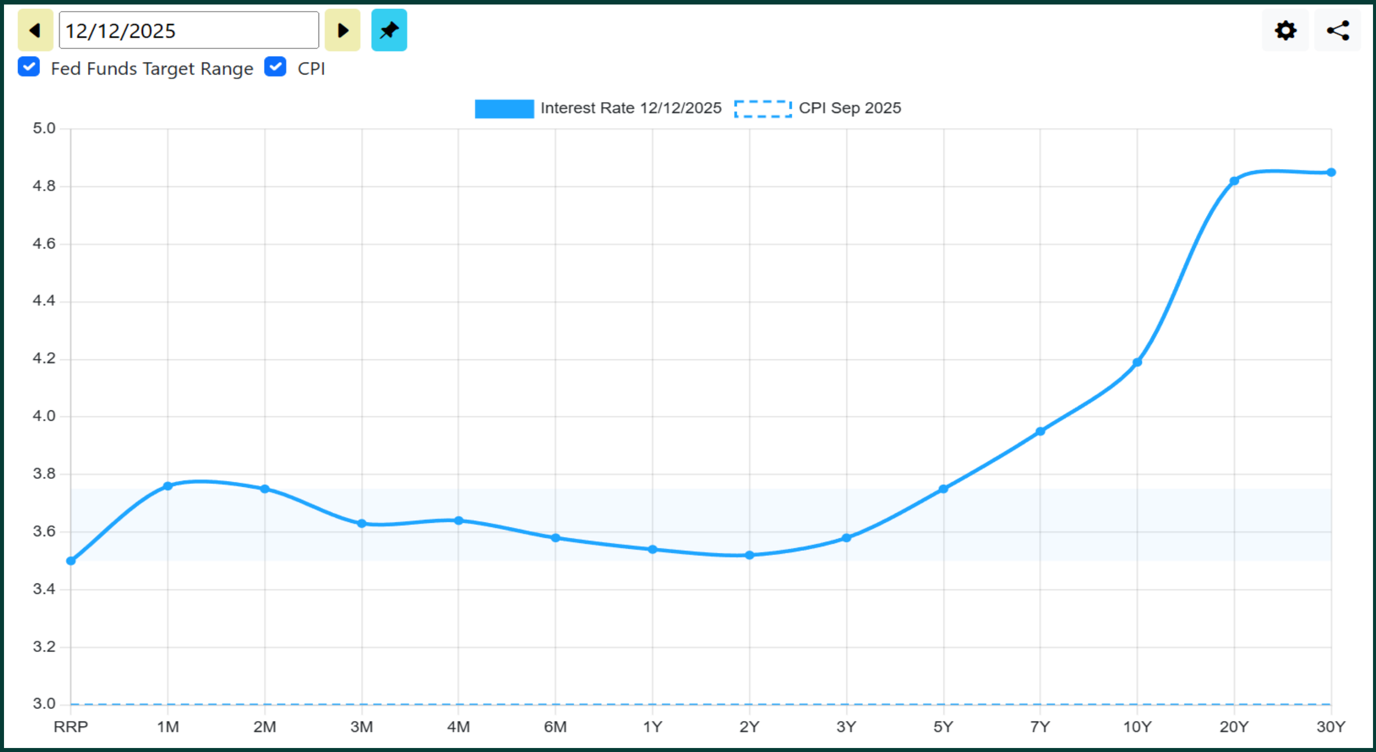

U.S. Treasury Yield Curve, *ustreasuryyieldcurve.com* (Treasury, Federal Reserve, BLS data), accessed December 12, 2025.

The current yield curve remains elevated and upward sloping beyond the front end. While near-term policy easing is being priced in, longer-term rates continue to reflect concerns around inflation persistence, fiscal supply, and consistently poor deficit health.

The Federal Reserve's recent decision to purchase securities at the lower end and belly of the curve is intended to reinforce transmission at the front end and stabilize short-term funding conditions. In the post-Covid environment, changes in policy rates have translated more slowly into the borrowing costs that matter most for households, including mortgages, auto loans, and consumer credit. This delayed translation should be expected through 2026 as well.

While easier policy may support asset prices and balance sheets at the top end of the distribution, consumer-facing businesses remain constrained by demand elasticity rather than financing conditions alone, suggesting that relief from rate cuts may be uneven and slower to materialize across the broader economy compared to previous easing policy periods before the pandemic.

Alliance Dependent Fragmentation

Global reorder is positioned to remain a top-level dynamic through 2026 that ripples through every facet of microstructure in domestic business. While fragmentation ultimately restructures relative strength, access to natural materials and energy production emerges as the infliction point of re-order dominance. Export controls and policy remain a central geopolitical tool - a reflection of earlier eras of disruptive technological advancement.

Against this top-level dynamic, there is still deep interdependence through 2026. The global semiconductor and supply chain ecosystem remains distributed. No single country can independently produce every critical component, however, restructured alliances can.

This reality is shaping the strategic calculus of both governments and multinational firms as they position for the next decade of technological competition.

U.S. - China relations remain a central risk factor for equities. Escalation with China over fundamental disagreements regarding Taiwan stands as the single most important geopolitical variable for the trajectory of U.S. artificial intelligence, related hardware manufacturers, technological progress more broadly, and, by extension, the domestic equity markets. The impacts of this dynamic are further complicated by China's deepening ties with key U.S. supply-chain diversification partners, most notably Vietnam, against the backdrop of an U.S. led push toward global fragmentation and reshoring.

China has managed to build as $1T USD trade surplus while most other nations struggle to grapple with U.S. tariff policy in a push towards fragmentation. In the first 11 months of 2025, China has established a surplus of $245B USD with South-East Asia, >28% increase in total surplus over the full year of 2024.[3] This sharp increase is almost fully attributable to a rise in transshipping engagement with the region to avoid U.S. tariffs. All the while, the aggressivity of U.S. tariff and export policy on China has only pushed China closer to acting for physical control of Taiwan.

3. Yahoo Finance. Jenny McCall – *How China racked up a $1T trade surplus*. December 13, 2025

Define a “Bubble”

1945 -1959 not 2000 - 2001

Labeling the AI boom a "bubble" by drawing parallels to the 2000–2001 dot-com collapse reflects a surface-level misunderstanding of both periods. The dot-com bubble was defined by speculative excess capacity built ahead of demand, while the AI expansion today is defined by persistent scarcity of capacity amid verified demand.

Two structural differences alone invalidate the comparison. First, the ability of AI to rewrite decades of globalization efforts into the formulation of fragile fragmented alliances, with a focus on full end to end AI supply chain development, over the past five years stands as a complete opposite of the momentum that characterized the late-1990s dot-com era. Second, equating AI with a bubble that collapsed under a bandwidth price vacuum and "dark" fiber build out on empty demand assumes commoditization. That analogy fails: fiber was homogeneous and incapable of compounding value, while AI systems generate parallel value through learning, reliability, and embedded knowledge, all characteristics that resist price collapse toward some marginal cost.

Today's AI infrastructure buildout reflects sustained demand across compute, power, and data centers. U.S. data-center vacancy rates are at historic lows, nearly all capacity under construction is pre-leased, and expansion is tied to identifiable workloads under long-term contracts. While both periods involved infrastructure built ahead of utilization, the dot-com era relied on speculative demand; AI investment is anchored to known counterparties and real revenue streams.

Corporate AI spending is largely funded through operating cash flows. Aggregate leverage among major firms remains low, underwriting standards have not meaningfully deteriorated, and available private credit exceeds $500 billion. Even as financing broadens, conditions remain far from historical speculative extremes.

Valuation dynamics further diverge. During the dot-com bubble, market values expanded far faster than earnings. Cisco's stock rose roughly 40x between 1995 and 2000, while earnings increased only about 8x. In contrast, leaders of the AI boom have consistently demonstrated earnings growth that outpaces valuation expansion. Nvidia's earnings have grown approximately 20x over the past five years, while its share price has increased closer to 14x.[2]

This distinction captures the core difference between the two periods. The defining constraint in today's AI expansion is not speculative excess, but physical and structural bottlenecks. Limited power availability, aging grid infrastructure, constrained compute hardware supply, and shortages of skilled labor relative to contracted demand indicate that the AI boom is governed by tangible capacity constraints, not speculative demand or sentiment-driven price inflation alone.

2. J.P. Morgan Asset Management – *Promise and Pressure*. Data as of December 17, 2025

The current AI boom maps well foundationally against the period of 1945–1959. Both periods represent foundational computing eras defined by scarcity and physical constraints rather than excess capacity built on speculative demand.

From 1945 through the late 1950s, computing advanced from ENIAC to the integrated circuit. Demand for computation was real but narrowly distributed, supply was constrained by hardware limitations and power requirements, and specialized labor and intellect was the fuel for scalability.

Today's AI expansion mirrors this dynamic. Although the supply today mostly supports civilian level demand, un-supplied demand is anchored by hyperscale's, governments, and large enterprises deploying identifiable workloads under long-term contracts. Like early computing, utilization follows capacity buildout and each marginal improvement in hardware unlocks step-function gains in capability rather than marginal price compression.

Critically, both periods are characterized by non-linear capability gains, not commoditization. Early computing did not collapse in value as machines became more powerful; instead, new classes of problems became solvable, expanding demand. AI follows the same pattern: efficiency gains are absorbed by increasingly complex models and applications, sustaining pricing power and capital investment.

In contrast to the dot-com era's overbuilt, homogeneous infrastructure, both 1945–1959 computing and today's AI boom are defined by constrained supply, strategic importance, and foundational buildout. A period defined by conditions that support long-duration investment cycles rather than speculative bubbles.

Certain company-level decisions warrant caution even at this early stage of the AI expansion, particularly for firms whose exposure is weighted more toward software-as-a-service than toward physical hardware. As financial conditions ease through 2026, the probability increases that the next phase of AI investment shifts from being primarily cash-flow funded toward greater reliance on debt, introducing a new layer of execution risk as timing becomes more significant.

Additional risks emerge around circular funding structures, raising the question of whether capacity could eventually outpace demand later in the cycle. That said, improvements in contract design, particularly around payback periods compared to instances in other technological advancement cycles, and the anchoring of these deals to private credit markets enhance project durability and near-term safety, while shifting a significant portion of risk away from public-market sentiment and valuation volatility.

Overall, given the cycle positioning, observable demand, and persistent physical constraints defining the AI buildout, the risk of underexposure to the winning end of a K-shaped economic environment is materially greater than the risk that a technological shift as foundational as autonomous computation forms a speculative bubble this early in its lifecycle.

Behavioral Biases

There are numerous historical periods of technological advancement that could be used as reference points for analyzing the AI investment cycle. Yet the comparison most frequently invoked is the dot com bubble. A technological disruption that is both the most recent in collective market memory and one that operated under a fundamentally different structural framework than the AI ecosystem today.

The tendency to equate the current AI boom with the dot com era is less a function of shared fundamentals and more a reflection of how human cognition processes technological uncertainty. Memory recall, pattern recognition, and narrative simplification favor familiar historical analogies, particularly those associated with highly salient losses, even when the underlying economic and capital structures differ materially.

Narratives questioning the sustainability of the AI boom increasingly center on valuation extremes. What is most visible; price appreciation, market concentration, and capital inflows, is frequently treated as most informative. In the process, attention is diverted away from the defining characteristic of the current cycle: AI investment remains constrained by tangible, physical capacity, including power availability, data center space, and specialized hardware.

Claims of “excess optimism” are frequently applied to a cycle that, to date, has been characterized by consistent earnings surprises and cash flow derived investment. Assertions that this environment will lead to indiscriminate overbuilding, an inability to absorb supply, and a subsequent market collapse implicitly assume a rapid transition to marginal cost economics like fiber bandwidth. Such assumptions overlook the structural reality that AI compute cannot be scaled or commoditized on the same timeline or under the same physical constraints as fiber cables, laid on empty demand, during the dot com bubble.

Part 2. Artificial Intelligence

An Evolution as Significant as the Rise of Computing Itself

Cycles

While corporate level adoption of Artificial Intelligence has been slower compared to that of the consumer, demand rivals full structural capacity and earnings continues to be the driving factor in growth –a stark contrast to past technological disruptions.

Today's Artificial Intelligence story resembles the rise of autonomous computing itself but from a cycle perspective the boom is far from representing a true paradigm shift. Heading into 2026, AI represents the earliest stages of a disruptive technological cycle, constrained by hardware, compute capacity, resources, and other tangibles.

Where do we position from here?

Capturing value from the Artificial Intelligence revolution becomes about positioning to benefit from those who can alleviate the physical constrains of capacity and compute capabilities. Who or what will be a vital player in allowing the Artificial Intelligence revolution to transform into a true Paradigm Shift?

Manufacturing Determines World Order, Technological Advancement Determines Manufacturing Dominance

In nearly every era since the Roman Empire, manufacturing superiority has determined the global order. Since the advent of autonomous computing, which broadly shapes our modern definition of "technology", the hierarchy of world manufacturing capacity has in turn been determined by leadership in technological advancement.

In the era of globalization, and the periods which shaped it, scale and timing of technological advancement alone has won and lost wars.

Over the past five years, artificial intelligence and the hardware underpinning its compute capabilities have radically expanded the frontier of what is possible in optimization related problem-solving and analysis. AI's impact on global re-ordering, its insatiable demand for compute infrastructure, and its measurable breakthroughs across fields such as economics, medicine, and materials science all signal the arrival of a new phase: one in which optimization-driven technologies become the decisive fulcrum for securing manufacturing superiority in the next era of history.

In just the past 12 months, Artificial Intelligence has cracked scientific, engineering, and economic optimization problems that were historically considered intractable.

Biology & Medicine: AlphaFold 3 (Google DeepMind + Isomorphic Labs) achieved unprecedented accuracy in predicting the structure and interaction dynamics of biomolecules, opening a new era of rational protein design.

Synthetic Biology & Materials: Meta's ESM3 created entirely new proteins with abstract functions including digesting plastics and forming bio-based industrial materials.

Materials Science: DeepMind's GNoME identified 2.2 million new crystalline materials, with 380k immediately synthesizable.

Energy Storage & Battery Chemistry: Collaboration between Microsoft and the Pacific Northwest National Laboratory produced a novel electrolyte formulation that resolves a long-standing bottleneck in both EV and grid-scale lithium metal batteries, improving stability and dramatically reducing cost.

Mathematics: DeepMind solved two long-standing conjectures, one in representation theory and another involving Kazhdan–Lusztig polynomials. Demonstrating AI's ability not just to accelerate applied sciences but to extend foundational mathematics itself.

Semiconductor Manufacturing: NVIDIA's ChipNeMo and CuLitho optimized core lithography workflows, reducing optical proximity correction computation time by orders of magnitude and helping preserve the viability of Moore's Law in an era of extreme design complexity.

In each instance, AI has compressed decades of classical-computing experimentation into months, redefining what is feasible in discovery and optimization. When considering how to position for the next paradigm shift, it is essential to recognize that these achievements, despite dramatically accelerating classical-computing workflows, remain fundamentally constrained by the limits of classical computing itself.

Total capital expenditure from the pioneers of Artificial Intelligence is projected to triple through 2026 relative to 2023 levels. Supporting this unprecedented surge, an estimated 73–93.5% of all new data-center capacity under construction is already pre-leased, while vacancy rates for existing facilities sit at an exceptionally low 2–4%.

Where earlier technological revolutions were bottlenecked primarily by intangible constraints such as computational speed, algorithmic efficiency, or chip performance, the defining challenge of the AI boom is unmistakably physical: the capacity of the underlying infrastructure to support a structural reordering of long-run supply and demand.

The United States enters this era with significant structural deficiencies. Grid modernization remains inadequate. Specialty materials production from grain oriented electrical steel to high-purity inputs for transformers and chipmaking lags strategic needs, large power transformer supply and efficiency is poor, and the nation faces a shortage of skilled technical labor required to deploy, maintain, and scale AI supportive systems. The next phase of AI leadership will not be determined solely by algorithmic superiority, but by a nation's ability to build, expand, and secure the physical backbone that enables AI to scale.

The Dual Dependency Constraint Loop and The External Third Variable.

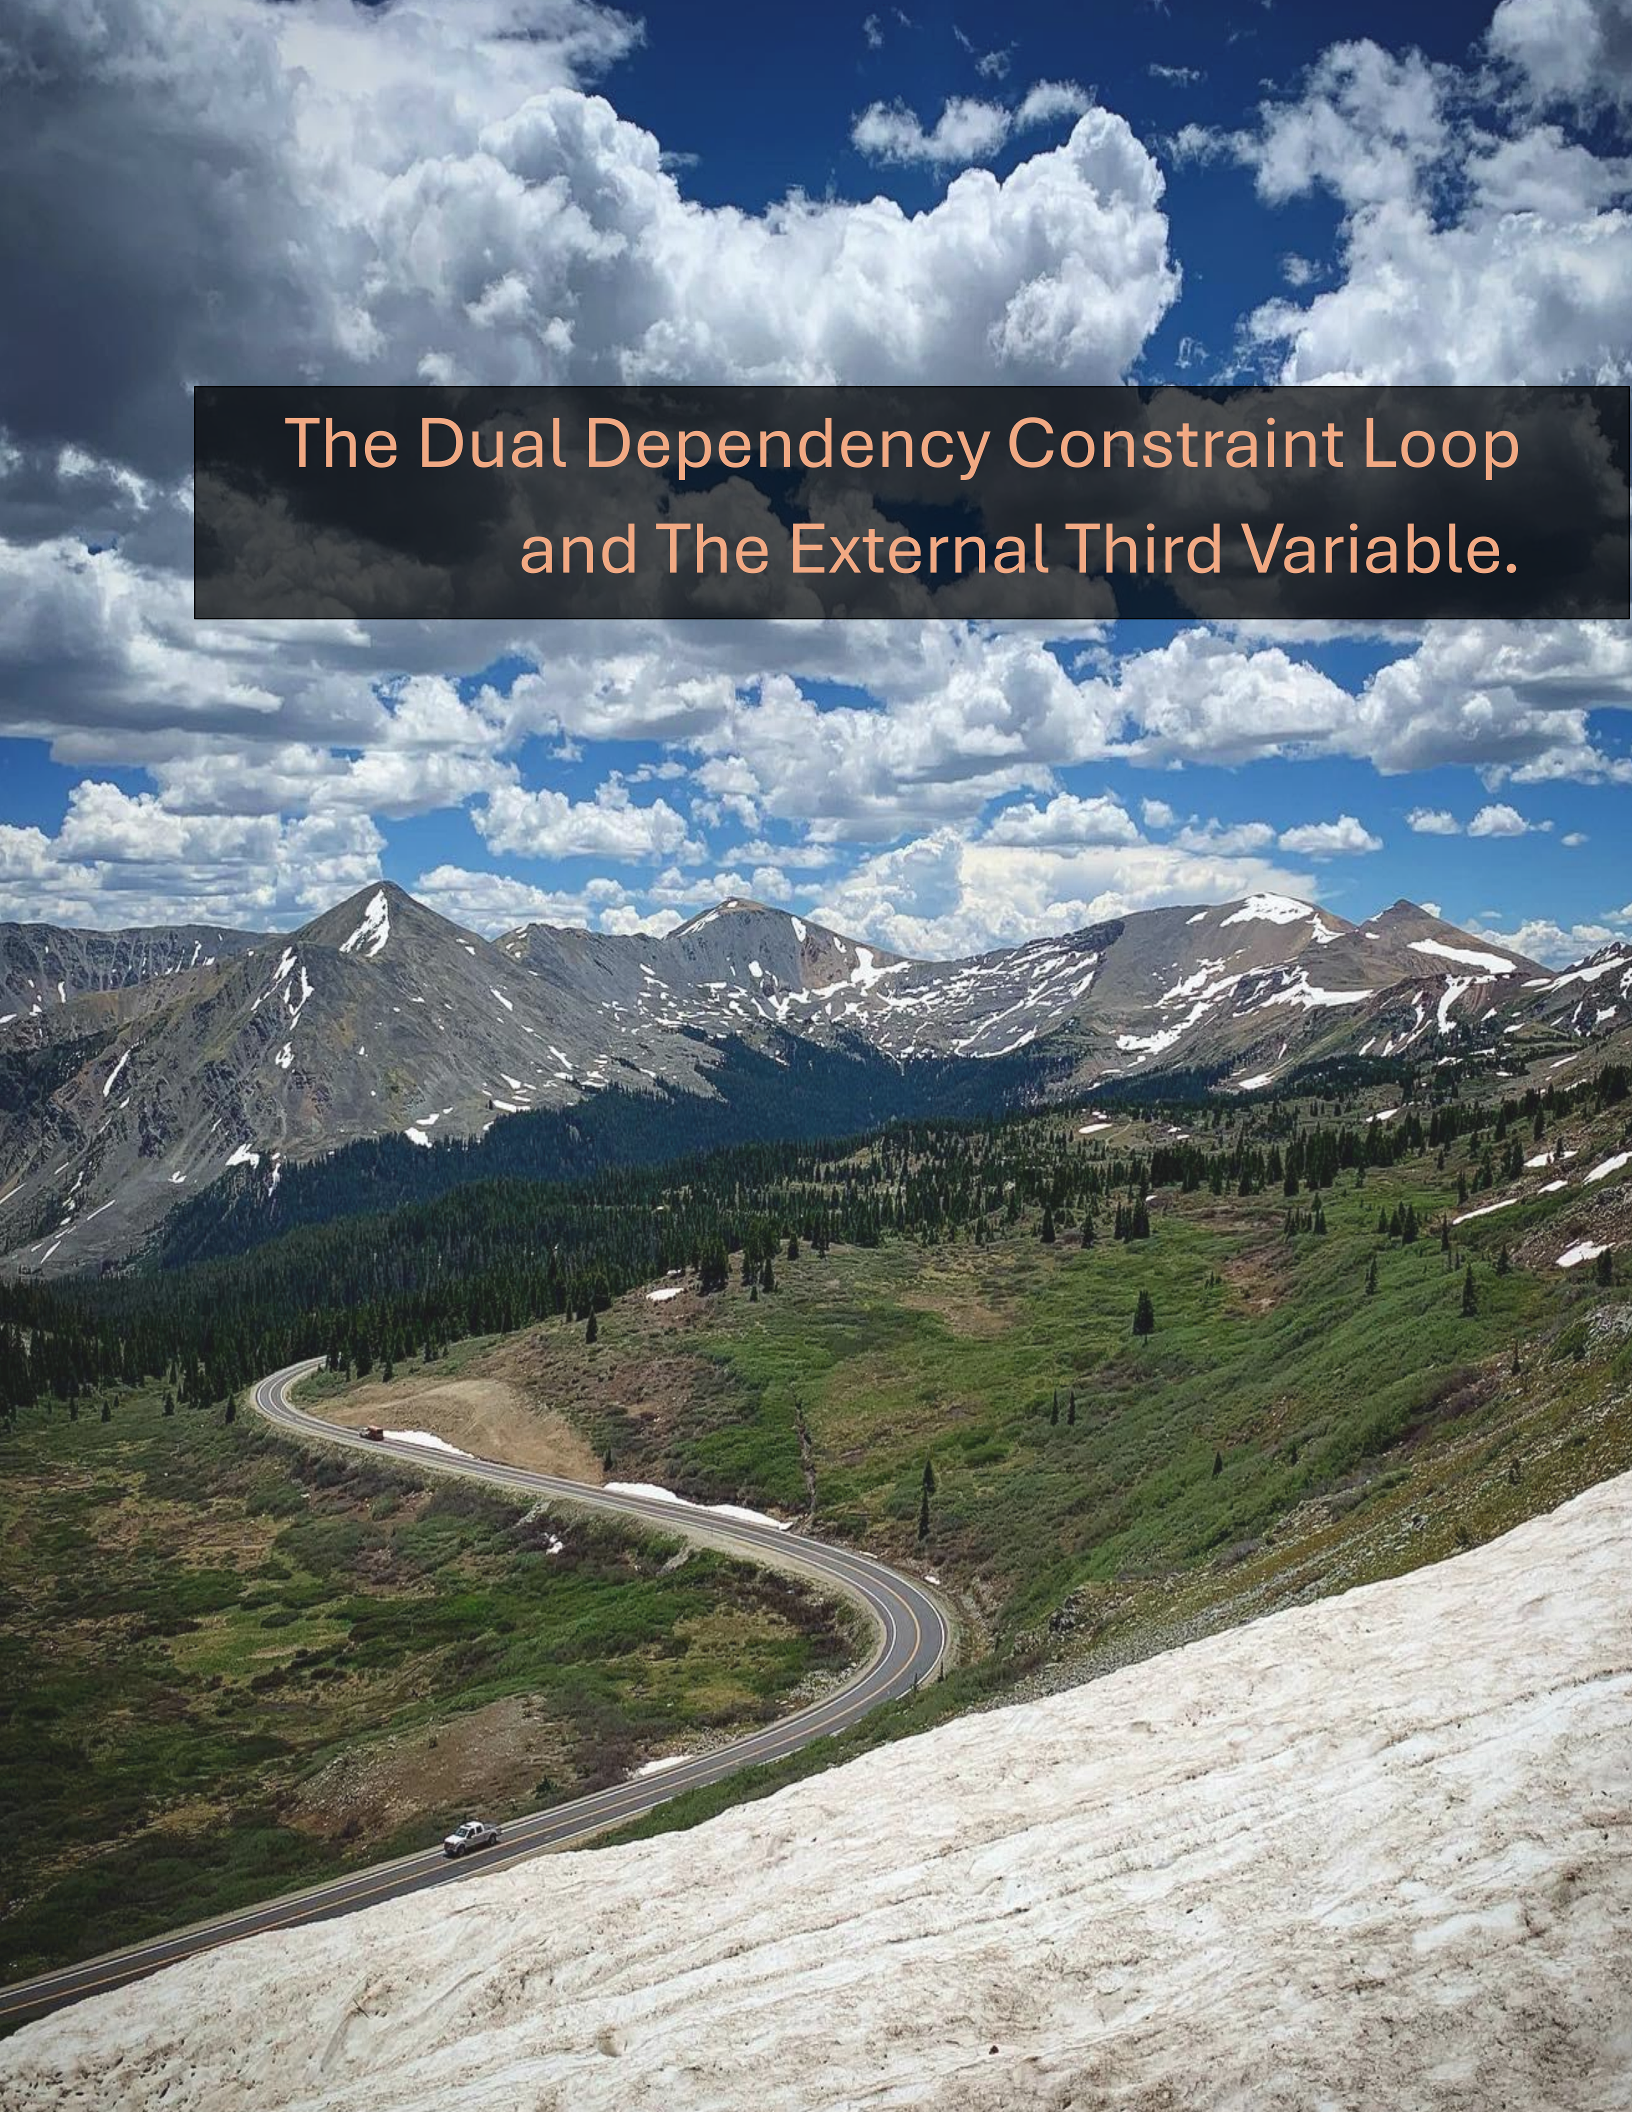

Self-reinforcing Constraints

Major bottlenecks in scaling artificial intelligence such as power generation, semiconductor capacity, advanced manufacturing, and the rejuvenation + scaling of electrical grids demands a dramatic expansion of labor input. Particularly in skilled trades, materials scientists, fabrication specialists, and construction workers.

Yet the United States, along with most advanced economies, is entering a prolonged period of labor scarcity: aging demographics, declining birth rates, chronically insufficient STEM pipelines, restrictive immigration relative to actual needs, and long training lead times for high-skill trades.

Duality Component 1: Labor scarcity implies that physical bottlenecks to scaling AI are constrained by the very input required to remove the bottleneck. AI cannot expand without massive physical infrastructure; power generation, transmission lines, semiconductor capacity, advanced manufacturing, and each of these requires substantial labor inputs.

Autonomation can help alleviate the impact on labor as a scarce input.

Duality Component 2: Autonomous systems can alleviate the labor gap by amplifying productivity and compressing development timelines. Yet these systems depend on the very infrastructure that is constrained by labor scarcity. Such systems require, existing compute capacity... existing semiconductor and fabrication capacity... existing power generation... existing transmission and grid buildout...

Duality Component 2 is dependent on Duality Component 1.

Automation can accelerate progress only after the physical foundations have been expanded, but expanding those foundations requires labor that is already scarce.

In an environment defined by a closed loop that is self-constrained by a scarce input, breaking the loop requires overturning the structural boundaries within which it operates. A system cannot resolve a constraint that originates from its own internal mechanics. An external variable that alters the parameters of the system itself is required to disrupt and ultimately solve a self-constrained cycle.

With Artificial Intelligence built on top of existing compute architectures and operating within a tangibly constrained environment shaped by scarce inputs, can one reasonably argue that AI itself constitutes a paradigm shift? Or is AI better understood as the foundation for recognizing that a new paradigm is required?

Thomas Kuhn's Paradigm Shift Theory

Thomas Kuhn argued in The Structure of Scientific Revolutions (1962) that scientific and technological progress can be broken down into an event structured timeline:

1. Normal Science: A period of stability where a community works within an accepted framework (a "paradigm") of shared assumptions, tools, and methods.

2. Accumulating Anomalies: Over time, enough problems emerge that cannot be solved under the existing paradigm, "anomalies" that reveal structural bottlenecks.

3. Crisis: The framework begins to break down. Existing methods and tools cannot efficiently scale with the demands placed on them.

Rober Noyce, "the Mayor of Silicon Valley." Co-Founder of Fairchild Semiconductor (1957) and later Intel (1968)

4. Paradigm Shift (Revolution): A fundamentally new framework replaces the old one. Not an incremental improvement but rather a complete redefinition of what is possible, changing the kind of solutions engineers can build.

5. New Normal Science: The new paradigm becomes the foundation for rapid progress that was impossible under the old system.

This structure maps remarkably well onto prior cycles of technological transformation such as the transition from vacuum-tube and lightbulb-based computing to silicon semiconductor chips. The transition to chip-based computing became normal science after a true paradigm shift and remains the operative baseline of modern computing today, still the foundation of Artificial Intelligence.

The Third Variable

A paradigm shift is not defined by “better tools”; it is defined by the emergence of an entirely new dimensionality in what tools can do. Artificial Intelligence remains rooted in a dimensionally defined sequential compute manner with parallelism achieved through advanced architecture rather than a natural law of the computational methods themselves.

Within the structure of Kuhn’s Theory, AI functions as Normal Science. AI is a powerful and accelerating refinement of the existing paradigm, that simultaneously exposes Accumulating Anomalies at a non-linear rate.

These anomalies arise because AI progress is bounded by the physical and structural constraints of classical computing and its demand of scarce resources: power limits, semiconductor scaling ceilings, material scarcity, grid constraints, labor shortages, and manufacturing bottlenecks. The more capable AI becomes, the more sharply it reveals the limits of the underlying paradigm that sustains it. Under Kuhn’s framework, AI cannot yet be defined as a true Paradigm Shift. It is the phenomenon that illuminates the need for one.

Drawing a parallel to the period between 1945 and 1959 (birth of electrical autonomous computation to birth of the efficient silicon semi-chip) the transition from Crisis to Paradigm Shift did not emerge from incremental refinement. It emerged from a new physical architecture that dissolved the constraints of the prior system. In the same way, the path from today’s AI-driven crisis of scaling to a true paradigm shift, and the only viable means of breaking the Dual Dependency Constraint Loop, is contingent on a variable external to the self-constraining loop.

The defining feature of nearly every observable bottleneck within the Dual Dependency Constraint Loop is its optimization complexity. Power allocation, grid design, chip fabrication, transmission expansion, labor substitution, material allocation, and manufacturing throughput are all optimization problems that grow exponentially with scale. Classical computing is fundamentally limited in solving these problems at the speeds required.

Therefore, the most plausible external variable - capable of altering the bounds of the system itself, is a shift toward adding dimensional quantity to computation. Many frontier computing technologies operate under different physical and mathematical laws. They unlock higher-dimensional optimization spaces that cannot be accessed by electron flow on classical silicon.

These technologies do not merely improve computation; they redefine what computation is. And in doing so, they represent the most credible foundation for a genuine Paradigm Shift capable of collapsing the structural bottlenecks that constrain AI today and breaking the Dual Dependency Constraint Loop that defines the current era.