The Dynamo and the Token

The AI buildout is the largest infrastructure bet in history. The fortunes won't go where everyone is looking.

by Alan Shimel

Editor-in-Chief, Founder and CEO, Techstrong Group

June 2026

Contents

— Prologue 4

01 Three Substrates: Electrons, Bits, Tokens 4

02 The Parabola 5

03 Why Intelligence Commoditizes—and Why It Isn’t Optional 7

04 What Electrification Actually Teaches 8

05 The Twist: Who Is the Utility, and Who Is GE? 10

06 Where the Value Migrates 11

07 A Disclosure, and a Bet 13

— References 14

Abstract

Every general-purpose technology of the last hundred and fifty years has been a fight over a single unit of account. Electrification gave us the electron. Computing gave us the bit. Generative AI is giving us the token. Each followed the same arc—heroic capex, hype peak, price collapse, and then a long second act lived almost entirely on the layer above. This paper traces that lineage and argues that the durable returns of the AI buildout will accrue not to the owners of the infrastructure, but to the builders who reorganize their work around abundant, cheap intelligence.

KEY CLAIMS

- Electrons, bits, and tokens are a lineage, not a series of analogies.
- Intelligence will commoditize. The bull case quietly assumes it.
- The fortunes don’t go to the grid—they go to what gets built on top.
- The integrated giants win as appliance makers who own a power plant.
- The neoclouds are the regulated utility of this era.

Prologue

There is a number making the rounds that most of us can't really get our heads around the enormity of. In 2026, the four largest hyperscalers—Alphabet, Amazon, Meta, and Microsoft—are on track to spend a combined ~$725 billion on capital expenditures[1], up roughly 77% from a 2025 figure that was itself a record. McKinsey projects that keeping pace with compute demand will require nearly $6.7 trillion in cumulative data-center investment worldwide by 2030[2], of which about $5.2 trillion is AI-specific. The International Energy Agency expects data-center electricity consumption to roughly double to around 945 terawatt-hours by 2030[3]—slightly more than Japan's entire annual consumption today—growing about 15% a year. Goldman Sachs sees U.S. data-center power demand more than doubling from 31 gigawatts in 2025 to 66 gigawatts in 2027[4].

These are not technology numbers. They are infrastructure numbers, the kind of growth the modern economy has produced only a handful of times: the railroads, the telephone/internet network, the interstate highways, and—the one that rhymes loudest—the electrification of America.

Everyone is watching the grid go up. Is just watching a mistake?

Not because the buildout isn't real or isn't consequential. It is the defining capital event of the decade, maybe the century. The mistake is assuming that the companies pouring trillions into generating and distributing intelligence are the ones who will capture the enduring wealth it creates. History is fairly emphatic on this point, and it says something uncomfortable for the prevailing narrative: when a foundational technology arrives, the durable fortunes rarely accrue to the owners of the core infrastructure. They accrue to the people who build on top of it once it becomes cheap, boring, and ubiquitous.

That is the argument of this essay. It is also, I'll admit upfront, an argument that runs against the grain of the loudest bull case in the market right now—and against the financial interests of a good number of the vendors who advertise across the network I run. I'll get to that. Let's start with the bull case, because it deserves a fair and full hearing. But first, a short detour through history—because the unit of account we are now arguing over is the third in a sequence, and the sequence itself is the most useful map we have.

01 Three Substrates: Electrons, Bits, Tokens

Every general-purpose technology of the last hundred and fifty years has, at bottom, been a fight over a single unit of account. Electrification gave the modern economy the electron, distributed as a kilowatt-hour. The computer revolution gave it the bit, distributed as a packet over a global network. Generative AI, this decade's defining buildout, is giving it the token—the atomic unit of machine cognition, billed by the million.

These three substrates are not metaphors for one another. They are a lineage. Each one became indispensable in roughly the same way: a heroic decade of capital formation, a wave of triumphalist coverage, a quieter decade of price collapse, and then—only then—a long, profitable second act lived almost entirely on the layer above. Henry Adams, who watched the first one happen in real time, stood in front of a forty-foot dynamo at the 1900 Paris Exposition and felt the ground move under him[5]. "The dynamo," he wrote, "became a symbol of infinity." He was right about the awe. He was wrong about where the money would end up.

Paul David, writing in 1990, gave us the canonical account of how electrons actually paid off[6]. Electric motors were available in American factories for roughly forty years before they showed up in the productivity statistics. The gains arrived only when factory owners stopped wiring dynamos into the bones of a steam-era plant and rebuilt the floor around the assumption of cheap, distributed, always-on power—the famous unit drive. The technology was necessary. The reorganization was sufficient. The fortunes followed the second one, not the first.

Bits ran the same play, faster. The 1990s telecom buildout laid millions of miles of fiber, lit it, and then watched the long-haul price of a bit collapse by something like four orders of magnitude. The companies that built the pipes—Global Crossing, Worldcom, the regional bells—mostly did not survive the decade after the buildout in any recognizable form. The trillion dollars of enterprise value created on top of cheap bits accrued to a different set of names entirely: Google, Amazon, Facebook, Netflix, and the long tail of SaaS companies who built their entire business model on the assumption that bandwidth was, for their purposes, free.

Tokens are now arriving on a similar schedule, only compressed. The capex wave is hitting today; the price-of-intelligence collapse is, on the most credible projections, a matter of two to four years; the reorganization wave—the unit-drive moment for cognitive work—has barely started. The rest of this essay is an argument about where, in that compressed sequence, the durable money is going to sit. It will help to keep the lineage in mind. Electrons, then bits, then tokens. Same story, three times. We are watching the third act open.

02 The Parabola

Daniel Newman, CEO of The Futurum Group, recently posted a chart that captures the optimist's case more cleanly than any argument could. Drawing on S&P Capital IQ historicals and Bloomberg consensus estimates (as visualized by the Financial Times), it plots the combined free cash flow of Alphabet, Meta, Microsoft, and Amazon from 2006 forward, with estimates carrying it out to 2030. The shape tells a story in three acts.

For years, free cash flow climbs steadily, peaking around 2024. Then, in 2026 and 2027, it gets crushed—the capex wave finally hits the income statement, and Amazon's line dips negative. The skeptics seize on this: all that spending, and cash flow is going backward. And then, in the estimate years, the line goes vertical. Roughly $210 billion in 2028. Around $410 billion in 2029. More than $630 billion by 2030.

Newman's reading is that this is simply how investment cycles work[7]. The leaders of these companies aren't deploying hundreds of billions this year, and trillions over the next several, without long-term returns in mind. The cash-flow dip is temporary—the cost of building the AI factories required to generate returns that surpass anything in their history. Underinvestment, in his framing, is a bigger risk than overinvestment. The returns go parabolic. "Strap in," he writes.

Figure 1. The parabola

Combined free cash flow of Alphabet, Amazon, Meta, and Microsoft, 2006–2030E (USD billions)

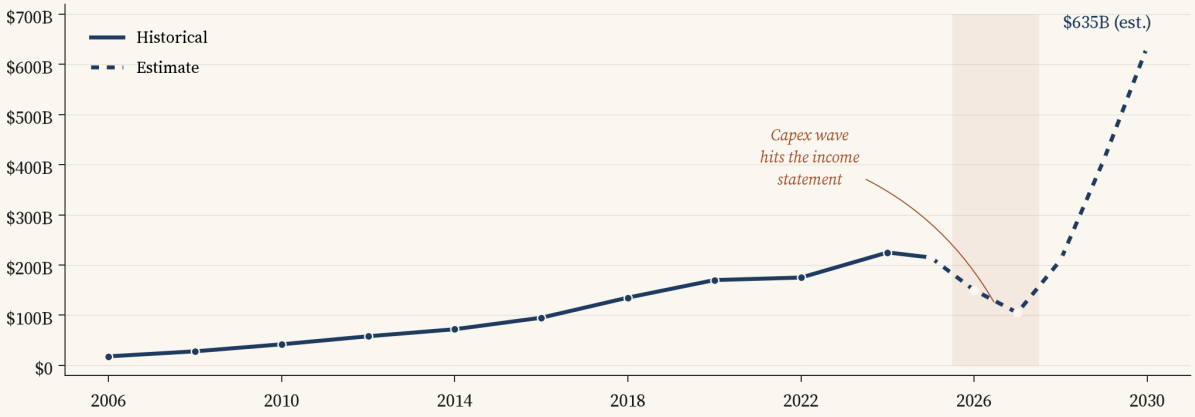

Source: S&P Capital IQ (historicals); Bloomberg consensus (estimates). Chart construction follows Daniel Newman / Financial Times.

Figure 1. Combined free cash flow of Alphabet, Amazon, Meta, and Microsoft, 2006–2030E. The 2026–27 trough is the capex wave hitting the income statement; the 2028–30 ramp is the optimist's case. Historicals via S&P Capital IQ; estimates via Bloomberg consensus.

He may be right. I want to be clear about that, because the easy move here would be to wave the chart away as hype, and it isn't hype—it's a serious claim backed by serious money, and the spending is, if anything, accelerating beyond these figures[8]. His own numbers actually run hotter than the consensus I cited above: more than $10 trillion in AI infrastructure by 2030, and by Futurum's estimates, more than $5 trillion in 2035 alone.

The chart cannot show you the thing that matters, though. A graph of cash flows captured during a buildout cannot tell you anything about returns retained after the technology it produces becomes a commodity. Those are two different questions, and the entire argument turns on the gap between them.

03 Why Intelligence Commoditizes—and Why It Isn't Optional

The frontier-model race looks, today, like a contest no one could ever win permanently. Each new release tops the last. The gaps between labs feel meaningful. Surely, the reasoning goes, whoever owns the smartest model owns the future.

I think that's a misread of where this is heading, and the reason isn't a forecast about model architectures. It's an economic constraint that the entire business case for AI quietly depends on.

AI displaces human labor only where it is cheaper than the human doing the task. That is the whole proposition—the reason a CFO signs the check. And right now, for a large and growing share of real-world workloads, the math is uncomfortably tight: once you account for the tokens consumed by reasoning, retries, context, and orchestration, running the work through a model can cost more than paying a person to do it. This is what all the industry talk of "tokenmaxing"—the scramble to wring more useful output from fewer tokens—is really about. It is evidence that the cost ceiling is real and already shaping behavior.

The implication: if token costs don't fall substantially, AI adoption stalls at precisely the margin where the value was supposed to come from. Which means there is relentless, structural, downward pressure on the price of intelligence—not because the labs are generous, but because their customers' ROI case collapses without it. Price compression isn't a possibility that might befall the model providers. It is a precondition the whole edifice requires in order to pay off.

This reframes Newman's parabola. That vertical line at the end of his chart assumes mass adoption. And mass adoption assumes intelligence gets cheap enough to clear the human-cost bar. The bull case and the commoditization case are not opposed. The parabola depends on the very price collapse that strips pricing power away from raw intelligence. The optimist's own math carries you to the commoditizer's conclusion.

There's a second force, subtler but just as corrosive to pricing power: substitutability. Different brands of gasoline make elaborate claims about additives and engine protection, and yet most drivers simply pull into whatever station is convenient and cheap. They "just get gas." Basic intelligence is heading the same way. I started writing the research behind this piece in one AI tool, moved to another, and may well finish in a third—and I lost almost nothing in the switching. When the cost of moving between providers approaches zero and the quality difference on routine work is marginal, the product is a commodity no matter how miraculous it looked three years ago.

And increasingly, the switching isn't even a human choice. The emerging default architecture routes each request to the cheapest model that passes the evaluation—provider-agnostic middleware that treats models as interchangeable inputs.

When indifference is engineered into the stack as a design principle, commoditization isn't a behavior you observe. It's a structural fact.

Intellectual integrity demands that we note this. A token is not a watt. Electrons are perfectly fungible; model outputs differ in quality, voice, and reliability, and for some high-value work—senior-engineer hours, large-scale legal review—AI is already dramatically cheaper than the human alternative, and there the cost ceiling doesn't bind at all. So intelligence won't commoditize into a pure commodity like raw electricity. It commoditizes into a differentiated commodity—more like gasoline with additives, or specialty coffee: thin margins on the base product, with room for some brand and switching-cost premium at the edges, and genuinely specialized models for specialized tasks. That nuance doesn't weaken the thesis. It is the thesis, stated precisely. The base layer—good-enough general reasoning, the bulk of the volume—commoditizes. The premium frontier holds a thinner and thinner slice. And the fortunes, as always, follow the volume.

04 What Electrification Actually Teaches

This is where the history stops being a flourish and starts doing real work.

When electric power arrived, it too looked like a business that whoever controlled it would own forever. Samuel Insull, the most important figure in the industry's early decades, argued that electricity was a natural monopoly—and, crucially, that monopoly obligations demanded public regulation[9]. He got his wish. The regulatory compact that followed[10] gave utilities protected service territories in exchange for universal service, cost recovery, and a "fair"—which is to say, capped—return. Formal state utility regulation in the U.S. is commonly traced to Wisconsin in 1907[11]; the model spread across the country.

That compact is the whole point. Once electricity became essential infrastructure, society stopped letting its owners price it like a scarce miracle and started treating it like a public good: affordable, reliable, universal. The owners of the grid did fine. They did not get rich in the way the technology's importance might suggest, because the moment their product became indispensable was the same moment they lost the right to extract its full value.

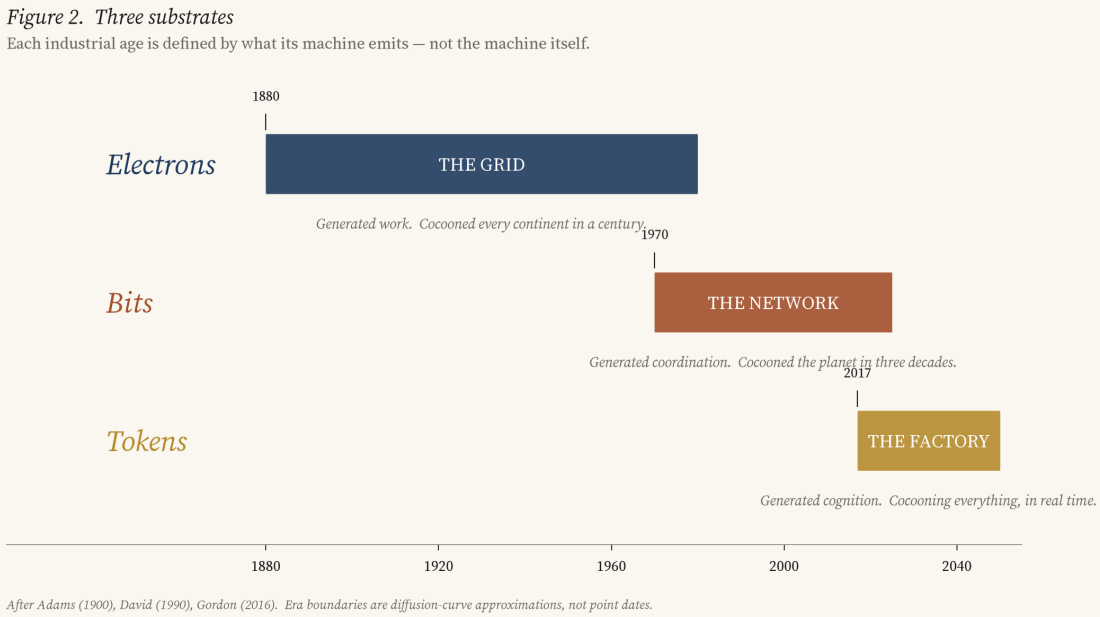

Figure 2. Three substrates, one pattern. Electrons (1880–), bits (1970–), and tokens (2017–) each followed the same arc: heroic capex, a hype peak, a price collapse, and then a long, profitable second act that lived on the layer above the grid.

So where did the fortunes go?

They went to General Electric[12]—which is instructive, because GE didn't only build the dynamos. Formed in the 1892 merger of Edison General Electric and Thomson-Houston, it built the generators and the appliances, the industrial systems, the motors, eventually aviation and healthcare and finance. It captured the grid layer and the build-on-top layer at once. Hold that thought; I'll come back to it in a moment.

And the fortunes also went, more broadly, to the manufacturers who reorganized their entire operations around the assumption of cheap, always-available power. This is the part most people get wrong about electrification, and the economist Paul David nailed it in his 1990 paper "The Dynamo and the Computer"[13]. Electric motors did not lift factory productivity for roughly forty years after they were available. The gains came only when factory owners stopped dropping electric motors into layouts designed for steam—a single central engine driving everything through a maze of shafts and belts—and rebuilt the factory floor around small, independent "unit drive" motors at each workstation. The technology was necessary but not sufficient. The payoff went to the reorganizers. The 1920s manufacturing productivity surge[14]—the boom that built the modern American economy, documented in Robert Gordon's economic history—was that reorganization finally cashing out, decades after the dynamo itself was old news. (The pattern is recurring: economists now describe a "productivity J-curve" for AI[15], in which measured gains lag the investment by years while firms quietly rebuild around the new capability.)

One more piece of the history is worth mentioning, because it complicates a too-clean story about markets. Private utilities largely declined to electrify rural America; the population was too sparse to be profitable. It took federal intervention—the Rural Electrification Administration in 1935[16], followed by the Rural Electrification Act of 1936—to finish the grid. When a technology becomes foundational enough, the market alone may not deliver universal access at socially acceptable prices. That is a live question for AI, and one I'll return to in future work.

05 The Twist: Who Is the Utility, and Who Is GE?

Now we can put the chart and the history together, and the picture sharpens into something more useful than "infrastructure owners lose."

Look again at the four companies in Newman's parabola. None of them is a pure grid operator. Each owns the intelligence infrastructure and a dominant application layer on top of it—Search and YouTube, Office and Copilot, Instagram's ad engine, AWS. They are not the regulated utilities of this story. They are GE: dynamos and appliances under one roof.

Figure 3. Who plays whom

Mapping the electrification era's two archetypes onto today's AI buildout.

The integrated giants own the grid and the appliance. The pure-play infrastructure sellers own only the grid.

Figure 3. The roles, then and now. The regulated utilities of the AI era are the pure-play infrastructure sellers; the integrated giants who own the grid and the appliances on top of it are this century's General Electric.

Which reframes the parabola entirely. If that vertical cash-flow line materializes, it almost certainly arrives through the application layer—ad optimization, cloud margin, Copilot seats, agentic products—not through selling raw tokens at a premium. The integrated giants may well win. But they'll win as appliance makers who happen to own a

power plant, not as power companies. Read that way, the bull chart quietly supports the up-the-stack thesis rather than refuting it.

So who plays the regulated utility in this drama—the essential, indispensable, margin-compressed grid operator? The pure-play infrastructure sellers. The neoclouds. The standalone data-center operators. The GPU-as-a-service platforms renting out undifferentiated compute. They own the wires without owning the appliances. They are building genuinely essential capacity, and for exactly that reason they are the most exposed to the commodity logic: indispensable, fungible, and priced accordingly. They are not in Newman's chart. That absence is the whole story.

There's a further reason the grid may stay a hard business rather than a harvest. Utilities never got to "build once, harvest forever"—the grid demanded perpetual reinvestment, which is part of why their returns stayed capped. AI infrastructure may be worse on this count. GPUs depreciate on roughly three-year cycles; staying competitive means continuous re-spend just to stand still. The question to put against the parabola is whether this is a build-once-and-harvest asset or a run-to-stand-still treadmill. Electrification suggests grids are treadmills.

06 Where the Value Migrates

This is where it matters most for you, because the whole conversation is fixated on the layer below you—and many haven't clocked yet that this is where the value is actually heading.

If intelligence becomes a cheap, ubiquitous, differentiated commodity, then competitive advantage stops living in access to a marginally smarter model and starts living in what you build on top of the assumption that capable intelligence is everywhere and nearly free. The grid getting cheap is not a threat to the people who build the appliances. It is the precondition for their entire economy. Electricity didn't transform anything as an expensive novelty in a few showcase mansions. It transformed everything once it got cheap enough that you'd design a whole production line—an entire business—around assuming it away as a free input.

So the practical question for anyone building right now is not "which model is winning this month." It's "what am I building that still has a moat when the model underneath me is a swappable commodity."

Here are a few places where that moat actually lives:

Proprietary data and the customer relationship. The model is rentable by everyone; your data and your customers are not. The durable asset is the thing the commodity layer can't replicate.

Workflow and the reorganization of work. This is Paul David's lesson, aimed straight at you. The fortunes won't go to the people who bolt a model onto an unchanged process. They'll go to the people who rebuild the process around the assumption of abundant

intelligence—who design the equivalent of unit drive instead of wiring a motor to a steam-era layout. That reorganization is hard, slow, and exactly where the surplus hides.

Figure 4. Where the value migrates

The AI stack — and the direction of durable returns.

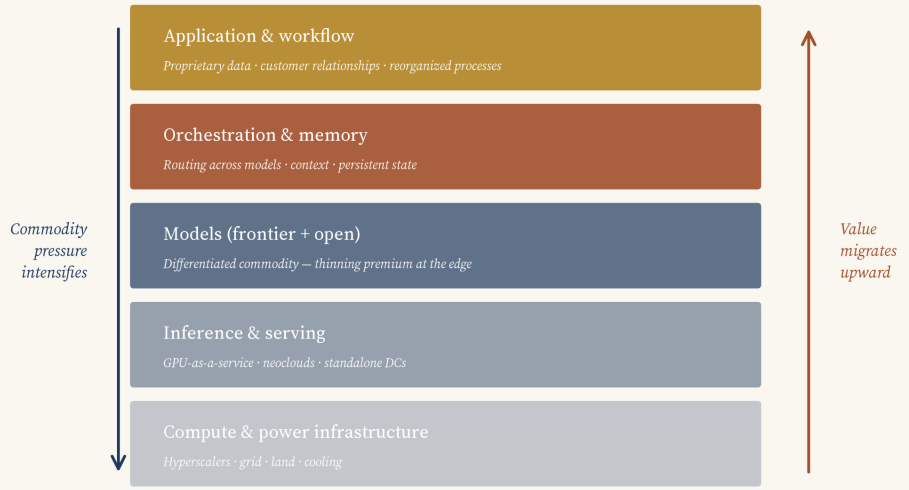

Figure 4. The AI value stack. Margin pressure intensifies downward as compute becomes a commodity; defensible value migrates upward into proprietary data, reorganized workflows, orchestration, and persistent memory.

Orchestration that assumes interchangeability. Build as though models are commodities—because they are becoming commodities—and the architecture that routes, evaluates, and composes across them becomes its own source of leverage.

Memory, context, and the layers that persist across whichever model is cheapest this quarter. The state that accumulates around the user is sticky in a way the model is not.

The most important word in that list is robust. You don't have to believe my timing to act on this. I think the rate of frontier improvement will slow and converge toward an equilibrium sooner than most expect—but you don't need that to be true on any particular schedule. Whether commoditization fully arrives in eighteen months or four years, the same bet pays off across the whole range: build on the assumption that intelligence becomes cheap, and don't stake your edge on privileged access to a model that won't stay privileged. The people who win are positioned for the world after the price collapse, not the world during the gold rush. The people who get rich in the gold rush are those selling picks and shovels. But that fortune fades once commoditization sets in.

07 A Disclosure, and a Bet

I feel compelled to note something here. I run a media network that depends on the very vendors this essay says are most exposed to margin compression—the cloud, infrastructure, and AI-tooling companies selling shovels in this gold rush. A thesis that predicts thinner returns for a chunk of my own advertisers is not a thesis I'd write if I were talking my book. I'm telling you anyway, because I think it's true, and because the only thing an analyst on an ad-supported platform actually has to offer is the willingness to say the inconvenient thing to the audience rather than for the sponsor.

As for the audience, for once, you're on the right side of the trade. The people who consume Techstrong content are practitioners—the builders. If this argument is correct, you are not spectators to the AI fortune. You are its likeliest beneficiaries, standing on the layer where value is migrating, provided you build for the world that cheap intelligence creates rather than the one expensive intelligence is leaving behind.

The capex spectacle will continue. The parabola may even arrive. Strap in, by all means. Just be clear about which seat you're in—because if a century of electrification is any guide, the gold rush won't end with the owners of the grid. It will end with the builders who make that grid indispensable, and then quietly capture everything that gets built on top of it.

That pattern—where value goes when a foundational technology becomes cheap—didn't begin with electricity, and it won't end with AI. The railroads, the telephone network, fiber optics, the personal computer: each tells a version of the same story, and each has something specific to teach about this moment. Those are the next essays. This was the first.

Figures are current as of mid-2026 and reflect the most recent published projections at the time of writing; capex and power-demand estimates are revised frequently.

References

1. Yahoo Finance, “Meta, Microsoft, Amazon, and Alphabet are about to spend a shocking amount of money to dominate the AI era” (2026). https://finance.yahoo.com/sectors/technology/article/meta-microsoft-amazon-and-alphabet-are-about-to-spend-a-shocking-amount-of-money-to-dominate-the-ai-era-115359575.html
2. McKinsey & Company, “The cost of compute: A $7 trillion race to scale data centers” (2025). https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers
3. International Energy Agency, “Energy and AI — Executive Summary” (2025). https://www.iea.org/reports/energy-and-ai/executive-summary
4. Goldman Sachs, “US Data Center Power Demand Projected to Double by 2027” (2026). https://www.goldmansachs.com/insights/articles/us-data-center-power-demand-projected-to-double-by-2027
5. Henry Adams, “The Dynamo and the Virgin,” ch. XXV of The Education of Henry Adams (1900/1918). https://www.gutenberg.org/files/2044/2044-h/2044-h.htm
6. Paul A. David, “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” American Economic Review 80(2), 1990, pp. 355–361. https://www.jstor.org/stable/2006600
7. Daniel Newman, summarized in Stocktwits, “Microsoft, Meta and Google just silenced AI spending critics” (2026). https://stocktwits.com/news-articles/markets/equity/microsoft-meta-and-google-just-silenced-ai-spending-critics-in-one-earnings-night-as-big-tech-capex-swells-to-725-b/cZBtCIgReEx
8. CNBC, “AI boom: Big Tech capital expenditures now seen topping $1 trillion in 2027” (2026). https://www.cnbc.com/2026/04/30/ai-boom-big-tech-capital-expenditures-now-seen-topping-1-trillion-in-2027-.html
9. Samuel Insull, “The Obligations of Monopoly Must Be Accepted” (1910), Yale University Energy History project. https://energyhistory.yale.edu/samuel-insull-the-obligations-of-monopoly-must-be-accepted-1910/
10. Yale Energy History, “Electricity and the Public Good: Private-Public Power Debates in the 1920s–30s.” https://energyhistory.yale.edu/electricity-and-the-public-good-private-public-power-debates-in-the-1920s-30s/
11. Forrest McDonald, “Samuel Insull and the Movement for State Utility Regulatory Commissions,” Business History Review (1958). https://www.cambridge.org/core/services/aop-cambridge-core/content/view/921F175ADA7D4939AE1BADB151811BB9/S0007680500005092a.pdf/samuel-insull-and-the-movement-for-state-utility-regulatory-commissions.pdf
12. Thomas A. Edison Papers, Rutgers University, “Edison General Electric Company.” https://edison.rutgers.edu/life-of-edison/companies/company-details/electric-light,-domestic/edison-general-electric-company
13. Paul A. David, “The Dynamo and the Computer” (1990) — see ref. 6. https://www.jstor.org/stable/2006600
14. Robert J. Gordon, The Rise and Fall of American Growth (Princeton University Press, 2016). https://press.princeton.edu/books/paperback/9780691175805/the-rise-and-fall-of-american-growth
15. Robert L. Bradley / AEI, “The Dynamo, the Computer, and ChatGPT: Explaining Today’s Productivity Paradox” (2023). https://www.aei.org/articles/the-dynamo-the-computer-and-chatgpt-explaining-todays-productivity-paradox/
16. EH.net Encyclopedia, “Rural Electrification Administration.” https://eh.net/encyclopedia/rural-electrification-administration/