THE THREE-DIMENSIONAL ECONOMICS OF AI

Rethinking returns when capital, revenue and technological velocity are endogenous

KRISH VASUDEVAN
Founder, TERAXYS LLC

TERAXYS LLC · SEPTEMBER 2026

The Three-Dimensional Economics of AI

Rethinking returns when capital, revenue and technological velocity are endogenous

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

PART I – The Claim on Capacity

Artificial intelligence is already becoming a material claimant on the world's productive capacity. Goldman Sachs estimates annual AI capital expenditure of about **$765 billion in 2026**, rising to **$1.6 trillion by 2031**.[1] For context, the OECD expects governments and corporations together to raise roughly **$29 trillion from debt markets in 2026**, and estimates that nine major AI players alone plan approximately **$4.1 trillion of capital expenditure between 2026 and 2030**.[2] The physical requirements are equally substantial. Global data centers consumed roughly **485 TWh of electricity in 2025**, with the IEA projecting consumption approaching **950 TWh by 2030**, just under 3% of global electricity demand.[3] Supplying that expansion requires generation, transmission, semiconductors, cooling, water, land and further capital. These forecasts already describe an enormous buildout, yet they necessarily impose a tractable path on technological development, hardware turnover and demand. If the frontier advances faster than those modeled trajectories, competitive reinvestment, power requirements and supporting infrastructure rise with it. **Under that trajectory, today's base case becomes tomorrow's floor.**

The emerging revenue evidence gives that investment a serious economic case. Anthropic's annualized revenue run rate exceeded **$65 billion by the end of July 2026**, rising from about $9 billion at the end of 2025.[4] Earlier financial projections indicated second-quarter revenue of at least $10.9 billion and approximately **$559 million of adjusted operating profit**, including model-training costs while excluding stock-based compensation.[5] The figures remain privately reported and the annualized run rate extrapolates a short period of current sales; their durability will become clearer as enterprise customers impose greater discipline on rapidly expanding AI budgets. Even with those qualifications, Anthropic provides substantial evidence that businesses will pay heavily for frontier capability, strengthening the economic argument for the proposed capital expenditure. **The question then becomes how that commercial success converts into durable economic returns once the cost of sustaining frontier capability is embedded in the calculation.**

PART II – The Red Queen

Current investment analysis already accommodates considerable uncertainty. Major investors assign forecast paths or sensitivities to revenue growth, margins, financing costs, infrastructure requirements, accelerator useful lives and technological obsolescence, then simulate a sophisticated range of outcomes from incremental capital deployment. Artificial intelligence presents a potentially different problem. Revenue, future capital requirements and the rate of technological advancement may

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

increasingly determine one another while changing at different rates over time. A higher economic return can attract greater investment; greater investment can accelerate the technological frontier; a faster frontier can shorten the economic life of existing capital and increase the amount required to maintain baseline competitiveness. **The variables used to calculate the return may therefore be dynamic, interdependent and partly endogenous to the return itself.**

The investment cycle increasingly resembles the **Red Queen effect**. Individual 2026 programs reach $220 billion at Amazon, $195–205 billion at Alphabet and $130–145 billion at Meta.[6] These commitments expand capacity while simultaneously establishing the competitive baseline for the next generation. A commercially superior model demonstrates that customers value greater capability, as Anthropic's recent growth suggests, strengthening the incentive for rivals to invest in closing the gap. Their response advances the frontier again, requiring the original leader to keep moving simply to preserve its position and the revenues supporting its existing capital base. A second dynamic resembles a **Chinese finger trap**: collective effort can tighten the constraint itself. Greater investment can accelerate technological progress, raise customer expectations and shorten the competitive life of earlier investment. **The industry is therefore running on a treadmill whose speed is partly determined by how hard its participants run—and whose gradient can steepen as they accelerate.**

That dynamic acts on capital through two converging clocks that are already moving at radically different speeds. Major hyperscalers carry server and network infrastructure over estimated useful lives approaching **six years**: Microsoft and Alphabet use six-year lives for significant server equipment, while Meta uses 5.5 years for most server and network assets.[7] Yet the intelligence those assets support can move several generations within months. OpenAI progressed from GPT-5.4 in March 2026 to GPT-5.5 in April and GPT-5.6 in July, with material capability gains across professional and agentic workloads.[8] Hardware advances on another, faster clock than its accounting life: Nvidia now operates on an annual architecture cadence, with successive generations delivering large improvements in throughput, energy efficiency and inference economics.[9] **Capital expected to produce economic benefits for five or six years is supporting products whose competitive frontier can reset within months, while newer hardware simultaneously improves the economics of serving that frontier.** Older compute can remain physically productive and commercially useful while becoming progressively less suited to the capability, cost and efficiency demanded at the frontier. **Its physical life can therefore substantially exceed its competitive**

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

economic life. The more consequential question is prospective: what rate of return should an investor require on the next dollar of AI capital when the capital itself may be committed for years while the model and hardware frontiers can reset in months?

Part III — TRACER™: Technological Rate-Adjusted Capital Expected Return

That question becomes difficult because the next dollar enters an economic system whose principal variables move together over time. Three matter immediately: **capital**, the money required to build and sustain frontier capability; **revenue or economic value**, the commercial return generated by that capability; and **technological velocity**, the speed at which the frontier advances. Each can alter the other two: capital can expand capability and revenue; revenue can attract further capital; and technological progress can increase economic value while shortening competitive asset lives and raising future capital requirements. **Figure 1 illustrates how the relationship among these three variables can change across successive technological states rather than remaining fixed through an investment period.** The relationship is therefore a feedback system rather than a one-directional investment chain. Its behavior also changes through time: if increasingly capable models contribute more effectively to the coding, research, testing and engineering of their successors, development cycles can compress further. **The next dollar would then face a technological environment changing faster than the one against which its return was originally estimated.**

For a prospective investor, historical return becomes a fallacious guide to the next marginal dollar when the technological frontier moves materially within the investment horizon. The relevant questions are prospective: what revenue can the next dollar generate, over what competitive economic life, and how much additional capital will be required during that period to preserve the capability supporting that revenue? An accelerating frontier can alter all three before the capital is fully deployed. Infrastructure financed and constructed against an eighteen-month capability horizon could enter service after that horizon has already been reset repeatedly. **The hurdle rate on the next dollar therefore cannot be inferred from yesterday's realized return. It must reflect the prospective capital requirement and economic life of the investment in the technological environment in which it will actually operate.**

We call the framework proposed to examine this problem TERAXYS TRACER™ — Technological Rate-Adjusted Capital Expected Return. TRACER begins one step earlier by asking whether the principal variables can be

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

assigned independent forecast paths at all. A useful model would need to recognize that:

- **Capital, revenue and technological velocity are interdependent.** Where these relationships are sufficiently stable, the model may produce a jointly determined expected return or range of returns.
- **Technological velocity has its own rate of change.** A frontier advancing at a stable pace creates different expected capital lives from one that is accelerating. **The model may therefore produce a return that changes materially with the rate of technological acceleration rather than remaining stable through the investment horizon.**
- **The system may be path-dependent.** Similar starting conditions can produce different investment outcomes depending on the sequence of model breakthroughs, competitive spending and customer adoption. **The model may therefore produce multiple internally consistent return paths rather than a single equilibrium result.**
- **The expected return may become non-estimable.** Cyclical behavior, sufficiently unstable relationships or technological change occurring faster than the relevant capital horizon may prevent the model from producing a reliable expected-return range. In a more extreme regime, the investment horizon itself may cease to remain sufficiently stable for a conventional expected return to be meaningfully defined. **The model's result may therefore be indeterminacy or a finding that the conventional return calculation has broken down under the prevailing technological conditions.**

Figure 2 expresses these relationships as a coupled dynamic system and summarizes the different classes of result the model may produce.[10] TRACER attempts to estimate the sustainable expected return on marginal AI capital as revenue, capital requirements and technological velocity evolve together over time. **If the system does not support a stable estimate, that result is itself economically meaningful.**

If estimating the required return increasingly demands continuous analysis of interacting variables moving at technological speed, will human analysis alone remain sufficient? And if AI becomes necessary to evaluate the economics of AI itself, how much of the resulting economic judgment are we prepared to delegate to it?

The Three-Dimensional Economics of AI
Rethinking returns when capital, revenue and technological velocity are endogenous

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

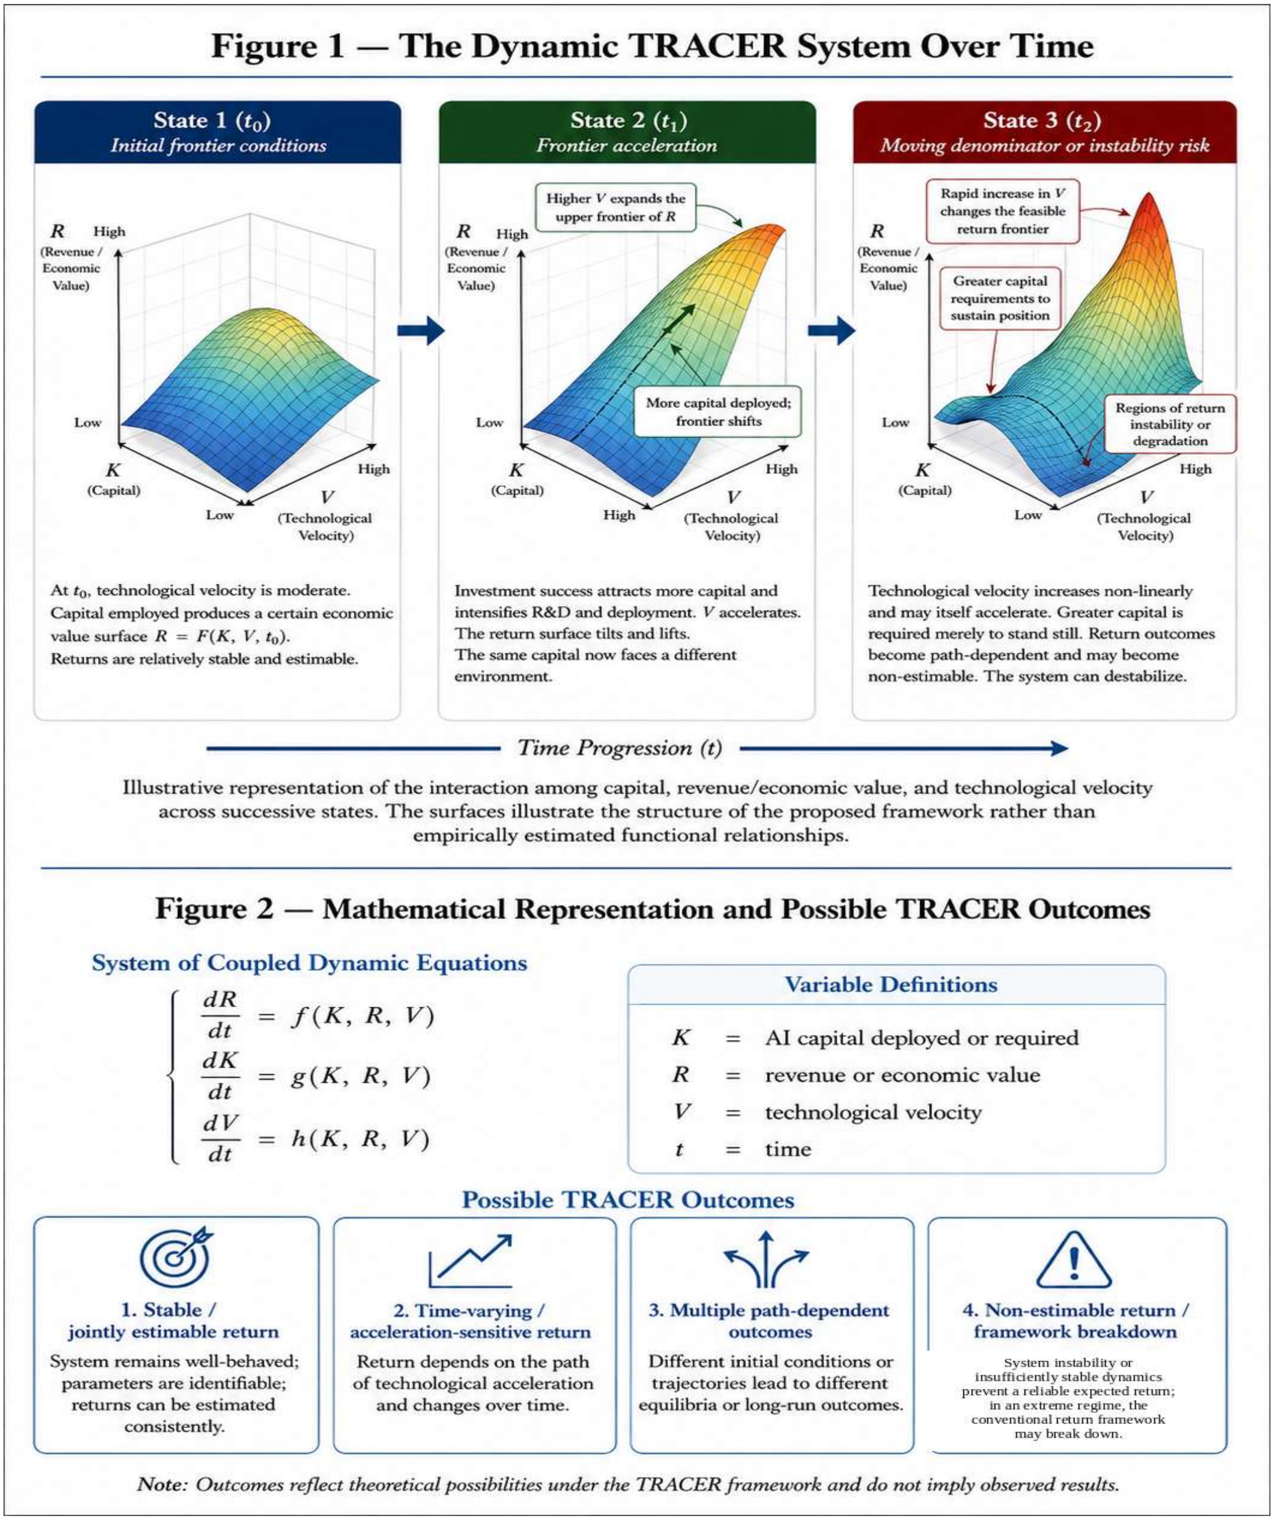

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

Part IV - The Real Economy

The same dynamics extend into the real economy. AI may generate productivity gains large enough to justify an extraordinary claim on scarce resources. Some impose a direct human opportunity cost: water, electricity and land diverted toward AI infrastructure are resources also required by households and communities. Others are more economically adjacent: capital, semiconductors and industrial capacity can be expanded, but directing them toward AI still affects financing costs, investment elsewhere and the timing of alternative production. Additional AI demand may itself induce new power generation, water infrastructure or semiconductor capacity and ultimately create greater aggregate value. The allocation problem therefore cannot be resolved simply by observing scarcity. It requires comparing the prospective productivity gain from AI with costs imposed elsewhere, across different people and different periods. If those trade-offs increasingly change at technological speed, do we eventually cross the Rubicon and allow AI itself to exercise judgment over the opportunity costs of its expansion? If we do, who determines the values embedded in that judgment?

The labor problem contains a harder contradiction. Evidence already supports the augmentation case: a 2026 study of more than 12,000 European firms estimated that AI adoption increased labor productivity by about **4% without reducing employment in the short run**.[11] Yet companies have also attributed roughly **112,000 corporate job cuts in 2026 to AI**, even though the causal attribution remains contested.[12] Augmentation therefore describes a possible transition rather than necessarily a stable employment equilibrium. If one AI-augmented professional can eventually produce the work previously requiring several, successful augmentation itself reduces the quantity of labor required for a given output. **The opposite outcome offers little protection: if the enormous capital commitment fails to generate the expected productivity and revenue gains, retrenchment across technology, infrastructure and the wider economy would create a more conventional employment shock.**

The macroeconomic issue is then compensation rather than headcount: high-income white-collar earnings support disproportionate consumption, taxation and household borrowing capacity. Governments may respond through retraining or a **Supporting Communities through Resilient Employment and Workforce Development initiative**. **But the proposition that income lost across multiple high-paying cognitive industries can be recreated elsewhere at comparable scale is increasingly suspect. Productivity can survive the worker whose income previously purchased its output.**

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

The distinction between firm-level success and economy-wide value is therefore critical. A company can improve margins by reducing labor while part of the resulting cost reappears elsewhere through unemployment benefits, retraining programs, income support and lower tax receipts. Revenue presents the mirror image. A dollar earned by an AI provider may represent genuinely new economic value, or spending displaced from software, consulting, search, coding, media, customer service or another AI provider. Aggregate output may still rise substantially. The harder measurement problem is the counterfactual: how much of that growth was created by AI rather than representing the trajectory the economy would have followed anyway? Econometric estimates can approximate that distinction, but the pre-AI economy cannot be rerun as a control case. A labor saving for one firm is income somewhere else; revenue captured by one platform may be displaced from another; and measured economic growth does not by itself reveal AI's incremental contribution. The return calculation ultimately requires all three to be distinguished.

Part V — The Measurement Problem

The TERAXYS TRACER™ framework proposes that the return on marginal AI capital may require a different measurement architecture when capital requirements, economic value and technological velocity evolve together at different and potentially unstable rates. It identifies the principal variables, their feedback relationships and the range of outcomes a dynamic model must accommodate: stable returns, materially changing returns, multiple path-dependent outcomes, an inability to estimate returns reliably, and—in an extreme regime—a breakdown of the conventional return calculation itself.

The next step is empirical. That work is naturally multidisciplinary. TRACER is offered as a framework for collaboration across mathematics, economics, technology and capital markets—and for a more fundamental question: before deciding whether AI is earning an adequate return, we may first need to determine what return means when the capital required to preserve it changes with the technology itself. If the variables governing that return ultimately become too dynamic for periodic human analysis, AI assistance may itself become necessary to solve the marginal-capital problem—requiring us to decide how much judgment we are prepared to delegate to the technology whose economics we are attempting to measure.

The Three-Dimensional Economics of AI

Rethinking returns when capital, revenue and technological velocity are endogenous

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

SOURCES AND NOTES:

1. *George Lee and Lucas Greenbaum, "Tracking Trillions: The Assumptions Shaping the Scale of the AI Build-Out," Goldman Sachs Global Institute, May 1, 2026. Goldman Sachs' baseline model estimates approximately $765 billion of annual AI capital expenditure in 2026, rising to $1.6 trillion in 2031, and approximately $7.6 trillion cumulatively between 2026 and 2031.*
2. *OECD, Global Debt Report 2026: Sustaining Debt Market Resilience Under Growing Pressure, OECD Publishing, Paris, 2026, DOI: 10.1787/e9d80efd-en. The OECD estimates that governments and corporations will raise approximately $29 trillion from debt markets in 2026 and that nine major hyperscalers are projected to incur approximately $4.1 trillion of cumulative capital expenditure between 2026 and 2030.*
3. *International Energy Agency, Key Questions on Energy and AI, IEA, Paris, 2026. The IEA estimates global data-center electricity consumption of approximately 485 TWh in 2025, rising to approximately 950 TWh in 2030, or around 3% of global electricity demand. Electricity consumption by AI-focused data centers is projected to triple over the same period.*
4. *Reuters, "Anthropic revenue run rate tops $65 billion, source says," August 17, 2026. Anthropic's annualized revenue run rate exceeded $65 billion by the end of July 2026, compared with approximately $9 billion at the end of 2025.*
5. *Reuters Breakingviews, "Anthropic's turbo-growth is only half the AI story," May 27, 2026. Anthropic projected second-quarter 2026 revenue of at least $10.9 billion and approximately $559 million of adjusted operating profit; the measure includes model-training costs and excludes stock-based compensation.*
6. *OECD, Global Debt Report 2026, 2026; Reuters, "Amazon lifts investment plans after strong cloud sales; shares jump," July 30, 2026; Reuters, "Google increases capex forecast again after cloud-driven quarterly beat," July 22, 2026; Meta Platforms, "Meta Reports Second Quarter 2026 Results," July 29, 2026. The OECD projects approximately $4.1 trillion of cumulative 2026–2030 capital expenditure across nine major hyperscalers. Current company guidance places 2026 capital expenditure at approximately $220 billion for Amazon, $195–205 billion for Alphabet and $130–145 billion for Meta.*

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

7. *Microsoft, Fiscal Year 2026 Second Quarter Earnings Conference Call, January 28, 2026; Alphabet, Investor Relations — Property and Equipment / Depreciation, 2026; Meta Platforms, 2025 Annual Report (Form 10-K), 2026. Microsoft describes relevant assets as depreciating over six years and discusses servers as having a six-year useful life; Alphabet generally depreciates servers and network equipment over six years; Meta extended the estimated useful lives of most servers and network assets to 5.5 years effective January 1, 2025.*
8. *OpenAI, "Introducing GPT-5.4," March 5, 2026; OpenAI, "Introducing GPT-5.5," April 23, 2026; OpenAI, "GPT-5.6: Frontier Intelligence That Scales With Your Ambition," July 9, 2026. The releases document successive improvements in professional, coding, agentic and other frontier-model capabilities over a period of approximately four months.*
9. *NVIDIA, "NVIDIA Kicks Off the Next Generation of AI With Rubin — Six New Chips, One Incredible AI Supercomputer," January 5, 2026; NVIDIA, "NVIDIA Vera Rubin Ramps Into Full Production to Power Agentic AI Factories Worldwide," May 31, 2026. NVIDIA describes an annual cadence for new generations of AI supercomputers and reports substantial generational improvements in throughput, efficiency and inference economics. Performance comparisons are vendor-reported benchmarks and may vary by workload and system configuration.*
10. *TERAXYS TRACER™ mathematical notation. In this conceptual representation, K denotes AI capital deployed or required to sustain competitive capability at a given point in time; R denotes revenue or economic value attributable to that capability; V denotes technological velocity, or the rate at which the relevant technological frontier is advancing; and t denotes time. Accordingly, dR/dt, dK/dt and dV/dt represent the rates of change in revenue, capital requirements and technological velocity respectively. The functions f, g and h represent the as-yet-undetermined relationships through which the three variables influence one another. The equations define the structure of the measurement problem rather than a solved mathematical model. No upper bounds on the variables are assumed in this conceptual representation; empirical development of TRACER would need to establish any constraints, boundary conditions or regime thresholds.*
11. *Iñaki Aldasoro, Leonardo Gambacorta, Rozalia Pal, Debora Revoltella, Christoph Weiss and Marcin Wolski, "AI Adoption, Productivity and Employment: Evidence from European Firms," EIB Working Paper 2026/02, European Investment Bank, January 13, 2026, DOI: 10.2867/1772538. Using matched data covering more than 12,000 non-financial firms, the authors estimate that AI adoption increased labor productivity by approximately 4%, driven by capital deepening, without an adverse firm-level employment effect in the short*

Krish Vasudevan
Founder, TERAXYS LLC
10.5281/zenodo.22033871

run; longer-term employment effects remain uncertain.

12. *Challenger, Gray & Christmas, "Challenger Report: Layoffs Fall, Hiring Picks Up; AI Leads for Fifth Straight Month," August 6, 2026; Reuters, "AI will not trigger employment collapse, staffing company Adecco Group says," July 23, 2026. Challenger reports that AI had been cited in 112,713 announced U.S. job cuts through July 2026, approximately 24% of all announced cuts. Employer attribution does not establish causation; Adecco CEO Denis Machuel argued that some companies may be attributing restructuring or underperformance-related layoffs to AI.*
13. *All illustrations and visualizations in this essay were created using artificial intelligence (AI) tools. The TERAXYS TRACER™ framework's mathematical representation was co-developed through iterative human–AI analysis. The author retains responsibility for the underlying thesis, assumptions, interpretation and conclusions.*