NBER WORKING PAPER SERIES

WHAT INVESTMENT DATA IMPLIES ABOUT THE AI TRANSITION

Jessica A. Wachter
Jonathan D. Wachter

Working Paper 35290
http://www.nber.org/papers/w35290

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2026, Revised October 2026

We thank Patrick Adams, Sylvain Catherine, Gary Gensler, Amy Huber, Alex Imas, Ioana Marinescu, Max Miller, Sean Myers, Stijn Van Nieuwerburgh, Fernando Stein, and seminar participants at Harvard and Wharton for helpful comments. We thank Frank Ma for valuable research assistance. The information, views, and opinions expressed herein are solely those of the authors and do not necessarily represent the views of Point72, its affiliates, or the National Bureau of Economic Research. Point72 and its affiliates are not responsible for, and did not verify for accuracy, any of the information contained herein.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Jessica A. Wachter and Jonathan D. Wachter. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What Investment Data Implies about the AI Transition
Jessica A. Wachter and Jonathan D. Wachter
NBER Working Paper No. 35290
June 2026, Revised October 2026
JEL No. E22, G12, O33

ABSTRACT

Five U.S. technology firms account for the majority of AI infrastructure investment, spending $404 billion (2022 dollars) in 2025, with nearly double that forecast for 2026. These firms risk bankruptcy unless expected profits grow commensurately. We embed this observation in a two-sector model with rare productivity booms. We calibrate the boom size to investment through 2026, implying that AI-sector productivity roughly tripled. We then add a two-year window with a 50% annual probability of another such boom. The implied GDP gain by 2030 ranges from 7% to 53%, with AI shares of GDP ranging from 8% to 36%.

Jessica A. Wachter
University of Pennsylvania
The Wharton School
Department of Finance
and NBER
jwachter@wharton.upenn.edu

Jonathan D. Wachter
Point 72
wachter@gmail.com

An interactive companion to this paper is available at
https://jwachter3.github.io/research/ai-transition/

1 Introduction

How much will artificial intelligence contribute to economic growth? Existing answers span an enormous range. Acemoglu (2024) applies a task-based framework and concludes that AI will raise total factor productivity by approximately 0.7 percentage points over ten years. At the other extreme, Amodei (2024, 2026) envisions industry-wide infrastructure spending reaching "multiple trillions a year by 2028 or 2029," with AI compressing decades of scientific progress into years.

Our contribution is to use investment data, combined with revealed preference, and a model for rare productivity booms, to estimate future GDP growth. Assuming the marginal investment is (conservatively) zero net present value implies a near-term increase in AI-sector productivity of a factor of roughly 3.3, a rare boom of the type studied by Tsai and Wachter (2016). We then use industry forecasts to 2030 to argue that firms are considering additional booms, but that these are far from certain.

The productivity implications from the current investment numbers are significantly greater than those of task-based analyses (Acemoglu, 2024, Babina et al., 2024, Yotzov et al., 2026), even those using more recent data. A second strand of the literature (Korinek and Suh, 2024, Restrepo, 2025, Trammell and Korinek, 2023) focuses on how wages might respond as human work becomes less necessary. If the physical and intelligence sectors are gross complements, AI-related growth eventually hits a wall (Kording and Marinescu, 2025). The extent to which these two sectors are gross complements is itself highly uncertain. Our model allows for long periods in which AI-related growth stalls. Chen (2026) focuses on the possibility of the singularity as a rare disaster. Rather than assuming the singularity, we back out the size of the first boom from investment data. The singularity is the scenario in which two further booms of this size occur. Van Nieuwerburgh (2026) examines how firms finance the AI buildout. We instead focus on the investment side. Our work complements

that of Eberly and Wang (2025), which studies the general equilibrium implications of an increase in productivity in one sector. To focus on the implications of the investment data, we assume an open economy, and therefore partial equilibrium.

Section 2 documents the facts. A small number of firms—Amazon, Alphabet, Microsoft, Meta, and Oracle—account for the vast majority of AI infrastructure spending. With five smaller firms that build AI capacity, their combined capital expenditure was $156 billion in 2022, rose to $255 billion in 2024, and is forecast at $782 billion for 2026, all in 2022 dollars. The 2022 and 2024 figures come from the firms' SEC filings and, for Microsoft, its earnings calls. The 2026 figure combines management forecasts for Amazon, Alphabet, Microsoft, and Meta with, for the other firms, capital expenditure in the first part of the year plus an extrapolation for the remainder. Section 2 reports these figures company by company and places them in the context of the aggregate economy.

Section 3 develops a two-sector model with rare productivity booms. The AI sector is exposed to the boom; the non-AI sector is not. The first boom entails an unanticipated jump in productivity, to which optimizing firms respond with a surge in investment. To justify investment through 2027, the boom size is roughly a multiple of 3.3 of productivity before 2024. This boom size informs the remaining calibration. With no other changes to productivity, this multiple alone implies a permanent increase in the level of GDP of approximately 5.6%.

However, industry forecasts, together with historical experience regarding other fundamental technological changes, suggest that the current productivity boom is not the end of the story. The rare boom model of Tsai and Wachter (2016) provides a means of capturing the skewness and optionality inherent in the future trajectory of AI. We model the occurrence of future advances as the outcome of a Bernoulli process with maximal uncertainty. We trace out a range of scenarios over 2028–2030, which we refer to as moderate, transformative, and singularity. The moderate scenario corresponds to the 5.6% gain mentioned above, already

an order of magnitude above task-based estimates. In the transformative and singularity scenarios, GDP in the post-window steady state is approximately 24% and 83% above its initial value.

If booms continue after 2030 at the largest annual probability at which firm values are finite, approximately 3.5%, expected cumulative GDP growth by 2050 is 32% in the moderate scenario (10th–90th percentile range 5.6–79%) and 363% in the singularity scenario (83–875%). Appendix D gives the implications of this uncertainty for the risk-free rate and the equity premium in an endowment version of the model.

2 AI capital expenditure

We use capital expenditure to measure the AI sector's investment: the capital added each year to its stock of data centers and computing equipment. Capital expenditure is not a defined term in either GAAP or IFRS. However, it is a number that firms report on earnings calls, and provide in earnings releases. Because it is not a required measure, firms report slightly different versions, though nearly always it is based on a line in the cash flow statement. Because we use the firms' own projections for 2026, our methodology requires using the firms' own measure, which we then back out for prior years. The goal is to come as close as possible to measuring additions to installed capital, in real terms. Using the firms' own definitions, despite the differences between firms, allows us to correctly capture the order of magnitude of the expenditures, though in some cases, our calculation may be conservative.[4]

Because AI infrastructure investment is highly concentrated, it is possible to approximate

[4]Specifically, by basing the calculation on cash flows, expenditure that is paid on credit may be delayed, or not show up at all. Leases may also not show up, unless the firm specifically defines capital expenditure to include them. Appendix A describes each firm's measure in detail, and discusses why these omissions are unlikely to make a substantial difference in the calculation.

the sector's investment from the regulatory filings and management forecasts of a relatively small number of firms. Table 1 reports annual capital expenditure for these firms, starting with the five largest, known as hyperscalers: Amazon, Alphabet, Microsoft, Meta, and Oracle. It includes five other firms that build AI capacity: xAI, the neocloud providers CoreWeave and IREN, and Applied Digital and Hut 8, which build data centers and lease them to AI tenants. To the extent that investment occurs outside of these firms, we undercount AI investment.[5] Dollar amounts in this paper are in 2022 dollars unless stated otherwise; the firm-level entries of Table 1 and the tables of Appendix A are the firms' own figures in current dollars.

Table 1 reveals that something unusual happened in 2024. Prior to 2024, the firms' aggregate capital expenditure was flat. Spending rose by two-thirds in 2024, and again in 2025, to $427 billion (the last row of Panel A). Capital expenditure in 2026 then implies an even greater increase – over 80%, to $782 billion, or approximately one trillion in 2026 dollars. We include 2027 to capture plans for continued expansion that the firms have already put in place: leases signed for data centers that open from 2026 on, commitments to buy equipment, and management's statements that spending will keep rising (see Appendix A.2). We take a conservative approach and assume that 2027 resembles 2026: each firm's 2027 capital expenditure keeps its net addition to the capital stock, measured in 2022 dollars, at its 2026 level, which is 2026 capital expenditure plus depreciation on the capital added in 2026. Table 1 shows this value in 2026 dollars and, for comparison, the values if we extrapolated recent growth.

Capital expenditure of the firms in Table 1 rose from 3.3% of total U.S. private fixed investment in 2022 to a projected 15.5% in 2026, and from 0.6% to 2.8% of GDP (Table 1, Panel B).

[5]Our identifying assumption is that capital expenditure is AI investment. The greater the amount of non-AI investment in 2022, the more conservative is this assumption.

Table 1: Capital expenditure of major AI infrastructure firms and its size relative to the U.S. economy

	2022	2023	2024	2025	2026	2027	
					Forecast	Assumed[a]	Projected[b]
Panel A. Capital expenditure ($ billions, current dollars)							
Amazon	58	48	78	128	200	229	388
Alphabet	31	32	53	91	185	218	400
Microsoft	28	41	76	118	175	202	300
Meta	32	28	39	72	138	161	248
Oracle	7	7	11	35	92	111	348
xAI (SpaceX AI segment)	—	—	6	13	55	68	—
CoreWeave	—	3	9	10	27	33	—
IREN	—	—	—	2	7	9	—
Applied Digital	—	—	1	1	6	7	—
Hut 8	—	—	—	—	2	3	—
Total	156	159	273	470	887	1,041	
Total in 2022 dollars[c]	156	153	255	427	782	917	
Panel B. Relative to the U.S. economy							
GDP ($T, current dollars)	26.1	27.8	29.4	30.9	∼31.8	∼32.7	
Investment ($T, current dollars)	4.7	5.0	5.3	5.6	∼5.7	∼5.9	
Share of GDP (%)	0.6	0.6	0.9	1.5	2.8	3.2	
Share of investment (%)	3.3	3.2	5.2	8.5	15.5	17.7	

Notes: Panel A. Values for 2022–2025 are actual capital expenditure, from the Forms 10-K of Amazon, Alphabet, Meta, CoreWeave, and Hut 8; the Forms 10-K and 10-Q of Oracle, IREN, and Applied Digital, with IREN's Form 6-K for its half-years; SpaceX's Form S-1 for xAI; and Microsoft's earnings calls. Values for 2026 are management's forecasts for Amazon (Form 8-K, February 5, 2026), Alphabet ($180–190 billion; Form 424B5, June 4, 2026), Microsoft (earnings call, July 29, 2026), and Meta ($130–145 billion; Form 8-K, July 29, 2026), with the midpoints of the ranges. For Oracle and Applied Digital, they are capital expenditure for December 2025–August 2026 from the Forms 10-K and 10-Q, with September–November equal to June–August; for xAI, CoreWeave, IREN, and Hut 8, they are capital expenditure in the first half of 2026 from the Forms 10-Q and 10-K plus twice that in the second quarter. Capital expenditure is each firm's own measure: cash paid for property and equipment for Alphabet, Oracle, xAI, CoreWeave, and Hut 8; the same, net of proceeds from asset sales and vendor incentives, for Amazon; the same plus principal payments on finance leases for Meta; equipment when received plus assets acquired under finance leases when the lease begins for Microsoft; cash paid for property, equipment, and computer hardware for IREN; and the capital expenditure of its AI segments for Applied Digital. The xAI entries are the capital expenditure of SpaceX's AI segment, and all of CoreWeave's capital expenditure is for AI. IREN's entries begin in 2025 and Hut 8's in 2026, when their spending turned from bitcoin mining to AI data centers; IREN's 2025 entry includes some mining hardware. "—" indicates less than $0.5 billion. Years are calendar years, except that Oracle's and Applied Digital's run from December of the previous year through November. [a]2026 capital expenditure plus depreciation on the capital added in 2026, in 2026 dollars, which keeps the firm's net addition to its capital stock, measured in 2022 dollars, at its 2026 level (Table A.5). [b]For the five largest firms, the year-over-year growth of the latest reported quarter, continued through 2027. [c]The totals converted to 2022 dollars with the consumer price index (CPI-U, annual averages; January–September for 2026); these are the values the calibration uses (Section 3.3). Appendix A gives the sources and calculations. Panel B. Investment is U.S. gross private fixed investment. GDP and investment for 2022–2025 are from the Bureau of Economic Analysis (NIPA Table 1.1.5, as published in September 2026; FRED series GDPA and FPIA), and the 2026–2027 values extend the 2025 figures at 3% nominal annual growth.

3 A two-sector model with rare booms

Our production framework builds on models developed by Gourio (2012) and Gomes et al. (2019), who study firms subject to rare disasters. We replace disasters with booms, and we break the assumption, introduced by Gourio (2012), that the shock moves productivity and the capital stock proportionally. In the disaster setting, the shock destroys both z and K in proportion, leaving the marginal product of capital unchanged; investment does not respond. Here, the boom raises productivity but leaves the physical capital stock intact, so that the marginal product of capital jumps upward and firms face a large gap between current and optimal capital.

3.1 Two sectors

The economy consists of two sectors: AI (a) and non-AI (n). Each sector $j \in \{a, n\}$ produces output using a Cobb-Douglas technology:

$$Y_{jt} = z_{jt}^{1-\alpha} K_{jt}^{\alpha}, \tag{1}$$

where α is the capital share, z_{jt} is sector-specific productivity, and K_{jt} is the sector's capital stock. Capital in each sector evolves according to

$$K_{j,t+1} = (1 - \delta_j) K_{jt} + I_{jt}, \quad j \in \{a, n\}, \tag{2}$$

where I_{jt} is gross investment and δ_j is the sector-specific depreciation rate. Assuming that each firm begins at its optimal level of capital before seeing the shock, the problem is scale invariant, and it suffices to consider a representative firm's optimization problem.

In each period, with probability p_t, a boom occurs: AI productivity jumps by a factor e^{ξ}, where $\xi > 0$ is the log change in productivity due to the boom. The non-AI sector is

unaffected. Formally:

$$z_{a,t+1} = \begin{cases} z_{at}\, e^{\xi} & \text{with probability } p_t, \\ z_{at} & \text{with probability } 1 - p_t. \end{cases} \tag{3}$$

For convenience, we assume neither sector is affected by normal-times volatility or by drift. We allow p_t to vary over time; in the calibration below, the variation takes the form of a structural break. Agents choose I_{jt}, and hence $K_{j,t+1}$, knowing z_{at} and p_t but not $z_{a,t+1}$.

3.2 Optimal investment

Firms in the AI sector maximize the present discounted value of output net of investment, assuming discount rate r_a; the non-AI sector's problem is analogous, with discount rate r_n.

$$V_{at} = \max_{\{I_{as}\}_{s \geq t}} \mathbb{E}_t \sum_{s=t}^{\infty} \frac{1}{(1+r_a)^{s-t}} \left[z_{as}^{1-\alpha}\, K_{as}^{\alpha} - I_{as} \right], \tag{4}$$

subject to $K_{a,s+1} = (1-\delta_a) K_{as} + I_{as}$ and the productivity process (3). Sectors face different required rates of return, r_a and r_n, reflecting differences in risk. Following an "open economy" interpretation, we treat these as exogenous, determined, for instance, by the required return of global investors willing to supply capital elastically at price r_j. Appendix D derives implications of rare booms for equilibrium rates of return.

The first-order condition with respect to $K_{a,t+1}$ equates the expected marginal product of capital to the user cost, the required return plus depreciation:

$$\mathbb{E}_t \left[\alpha \left(\frac{z_{a,t+1}}{K_{a,t+1}} \right)^{1-\alpha} \right] = r_a + \delta_a. \tag{5}$$

Evaluating the expectation over the boom:

$$\alpha \left(\frac{z_{at}}{K_{a,t+1}} \right)^{1-\alpha} \underbrace{\left[(1 - p_t) + p_t \, e^{(1-\alpha)\xi} \right]}_{\equiv \Omega(p_t)} = r_a + \delta_a, \tag{6}$$

where we define the boom option value:

$$\Omega(p) \equiv (1 - p) + p \, e^{(1-\alpha)\xi}. \tag{7}$$

Note that $\Omega(p) \geq 1$, with equality when $p = 0$.

Solving (6) for optimal capital:

$$K^*_{a,t+1} = z_{at} \left(\frac{\alpha \, \Omega(p_t)}{r_a + \delta_a} \right)^{1/(1-\alpha)}. \tag{8}$$

The optimal capital stock is proportional to current productivity z_{at} and increasing in the boom option value $\Omega(p_t)$. The non-AI sector satisfies the analogous condition with $\Omega = 1$ (no boom exposure):

$$K^*_{n,t+1} = z_{nt} \left(\frac{\alpha}{r_n + \delta_n} \right)^{1/(1-\alpha)}. \tag{9}$$

Proposition 3.1 (Capital ratio)**.** *Suppose a boom occurs at time t, so that $z_{at} = z_{a,t-1} \, e^{\xi}$, and that the boom probability shifts from p_{t-1} to p_t. Then*

$$\frac{K^*_{a,t+1}}{K^*_{at}} = e^{\xi} \left(\frac{\Omega(p_t)}{\Omega(p_{t-1})} \right)^{1/(1-\alpha)}. \tag{10}$$

Under the open-economy assumption, $r_a + \delta_a$ cancels from the capital ratio, which depends only on the boom size ξ and the shift in p. We can therefore identify ξ from investment data without taking a stand on the required return.

How a boom affects sectoral output. Because $K_{a,t+1}$ is chosen at time t, capital is predetermined when the boom is realized at $t+1$, and output responds in two stages. In the impact period, productivity jumps to $z_{at}\, e^{\xi}$ but capital is fixed at $K^*_{a,t+1}$, so output rises by the factor $e^{(1-\alpha)\xi}$:

$$Y_{a,t+1} = e^{(1-\alpha)\xi}\, z_{at}^{1-\alpha}\, (K^*_{a,t+1})^{\alpha}.$$

With an unchanged boom probability, firms then set $K^*_{a,t+2} = e^{\xi}\, K^*_{a,t+1}$, and if no further boom occurs, output rises by the remaining factor $e^{\alpha\xi}$ as capital adjusts:

$$Y_{a,t+2} = e^{\xi}\, z_{at}^{1-\alpha}\, (K^*_{a,t+1})^{\alpha}.$$

3.3 Calibration

All rates are annual. The capital share α is set to 1/3. The AI-capital depreciation rate, $\delta_a = 0.25$, reflects the shorter economic life of computing equipment versus general infrastructure. The non-AI depreciation rate is $\delta_n = 0.07$ (Cooley and Prescott, 1995). We assume a required real return on AI-sector capital r_a of 8%, a 2% real risk-free rate plus a 6% equity premium, slightly above what a rare-disaster model produces for the required return on the market (Wachter, 2013). Industry discount rates for AI investment are higher: Vanguard Group (2025) discounts expected AI earnings at a rate of 15%. Appendix C reports results for this higher rate, as well as for lower depreciation of AI infrastructure.

Before the boom, we assume that investment replaces depreciation: $I_j = \delta_j K_j$. The initial AI capital stock is then $K_{a0} = I_{a0}/\delta_a = 156/0.25 = \$624\text{B}$, where \$156B is the 2022 capital expenditure in Table 1, and the initial non-AI capital stock is $K_{n0} = (4{,}700 - 156)/0.07 \approx \$65{,}000\text{B}$, from U.S. gross private fixed investment of \$4.7 trillion in 2022 (Table 1, Panel B). We set the initial productivity level z_{a0} so that K_{a0} satisfies the optimality condition (8), which implies $Y_{a0} = (r_a + \delta_a)K_{a0}/\alpha$. Non-AI output is $Y_{n0} = (r_n + \delta_n)K_{n0}/\alpha$. We set the

required non-AI return r_n so that initial GDP, $Y_{a0} + Y_{n0}$, equals measured 2022 GDP, which gives $r_n = 6.1\%$.

2024–2027. Analogously to the rare-disasters approach of Barro (2006), we simplify the calibration by collapsing a multi-period event into an instant of time. We interpret the investment in Table 1 after 2023 as the firms' response to a productivity increase that occurred prior to 2024. We assume the boom probability is unchanged at the time of this first boom and through the 2027 investment decision, $p_t = p_{t-1}$, so that the Ω terms in (10) cancel and Proposition 3.1 gives $K^*/K_{a0} = e^{\xi}$. The frictionless model closes the gap between K_{a0} and K^* in one year; in the calibration totals, investment rises in each year from 2023 to 2027. We therefore calibrate ξ to the level of the capital stock rather than to the timing of investment: we assume that capital reaches its target by the end of 2027, so that the capital stock entering 2028 equals $K^* = K_{a0}e^{\xi}$.

The calibration totals are those of Table 1 in 2022 dollars, the last row of Panel A. Starting from K_{a0} entering 2023, the law of motion (2) gives

$$K_{a,2028} = (1-\delta_a)^5 K_{a0} + \sum_{t=2023}^{2027} (1-\delta_a)^{2027-t} I_{at}, \tag{11}$$

and our revealed-preference estimate of ξ equals

$$e^{\hat{\xi}} = \frac{K_{a,2028}}{K_{a0}}. \tag{12}$$

With these totals, $K_{a,2028} = \$2{,}048\text{B}$, so

$$e^{\hat{\xi}} = \frac{2{,}048}{624} \approx 3.28, \qquad \hat{\xi} \approx 1.19.$$

This is, approximately, the cumulative increase in productivity required to justify this level

of investment, assuming unchanged discount rates and zero NPV at the margin. The calculations that follow use the unrounded value of $\hat{\xi}$.

The previous calculations model a boom in AI productivity that has already occurred, and back out its size from the investment response. However, one characteristic of AI is the potential for AI itself to create further productivity changes, and thus to increase not only the level of GDP but its growth. This is not unprecedented, as Hanson (2000) documents how the industrial revolution led to long-run growth rates that would have been unthinkable in the centuries prior.

Modeling a very high long-run growth rate implies that the optimization problem that determines our equations is not well-defined. Moreover, given the free flow of capital, it is possible that growth could be realized in the near term, and then fall to some level that is higher than the current measured level, but nonetheless below a discount rate of 8%. As an illustrative calculation, we assume a window of time during which uncertainty is maximal, and after which probabilities fall to a steady-state level. This calculation captures the most striking feature of the AI transition, its uncertainty. Assuming a two-year window may seem restrictive. However, note that at some point the probability must fall, and the randomness in the length of the window will be captured by Poisson uncertainty. For example, if there is a 4% chance of a boom occuring after the window, this is equivalent to an 8% chance, that the 50% probability extends. The 90% and 10% confidence intervals shown below capture the range of uncertainty that include further booms beyond the 3 of the most optimistic scenario, as well as no booms beyond the first.

2028–2030. We consider that over a short period of time, p is at the level corresponding to maximal uncertainty, that is, $p = 0.5$. We assume:

$$p_t = \begin{cases} p_0 \approx 0 & \text{for } t < 2028, \\ p_{\max} & \text{for } t \in \{2028, 2029\}, \\ p_* & \text{for } t > 2029. \end{cases} \tag{13}$$

We assume that the rise in the boom probability in 2028, (13), is unanticipated through 2027. This allows us to calibrate ξ to investment through 2027, independently of $p_{\max}$. It captures the part of the boom that is already known: investment through 2026, with net additions to capital in 2027 held at their 2026 level. If instead the elevated probability already applied to the 2027 investment decision, so that $p_{2027} = p_{\max}$, the option-value term in Proposition 3.1 would account for part of the capital in place by 2028, and the same capital stock would imply $e^{\hat{\xi}} = 2.14$ rather than 3.28. We assume $p_0 \approx 0$ for convenience; a low value is realistic, as very few foresaw the growth potential of AI prior to 2024. At the core of our calculation is the idea that a boom is followed by an increase in probability.

With two independent draws at $p_{\max} = 0.50$,[6] governing booms in 2029 and 2030, there are approximately three outcomes. In the *moderate* scenario, which has probability $(1 - 0.5)^2 = 0.25$, no further booms occur during the window. In the *transformative* scenario, with probability $2 \times 0.5 \times 0.5 = 0.50$, exactly one additional boom arrives during the window: in 2029, which we call *transformative (a)*, or in 2030, *transformative (b)*. In the *singularity* scenario, with probability $0.5^2 = 0.25$, two additional booms occur, compounding the productivity gains.

[6] A Bernoulli random variable has maximal variance at $p = 0.50$. This assumption imposes the least informative prior on whether an additional boom occurs in any given year of the window.

Investment paths during the window. Write $K^*(z, p)$ for the optimal capital stock (8) at productivity z and boom probability p, and z_1 for productivity after the first boom. Each further boom multiplies z by e^{ξ}.

2028. When p rises to $p_{\max}$, the optimal capital stock entering 2029 rises from $K^*(z_1, 0)$, the capital stock entering 2028, to $K^*(z_1, p_{\max})$. The rise is the option-value term of Proposition 3.1. Investment in 2028 is the same in all four scenarios because the 2029 boom has not yet been observed.

2029. If no boom occurred (moderate and transformative (b)), investment just replaces depreciation. If a boom occurred (transformative (a) and singularity), the capital stock must adjust to the higher productivity, and investment is an order of magnitude larger.

2030. The boom probability reverts to p_*. Because $\Omega(p_*)$ is substantially below $\Omega(p_{\max})$, the optimal capital stock falls if no boom occurs; indeed, without a boom in 2030 (moderate and transformative (a)) it falls below what remains of the 2029 stock after depreciation, so investment without frictions would be negative. A boom in 2030 more than offsets the lower probability. This is one place where the paths are sensitive to the assumed probability of future booms: if agents anticipated further booms at a 50% probability, the capital stock would be maintained, and would then decrease or increase depending on whether such booms occurred. From 2031 on, absent further booms, investment replaces depreciation; if a boom occurs, investment catches up.

Partial adjustment. While this describes a frictionless investment sequence, in reality there are substantial frictions. A model of adjustment costs with the quadratic form (Hayashi, 1982) would predict that investment after a single shock follows a decelerating path regardless of the adjustment speed, whereas the data show acceleration. Ideally, one would model adjustment costs that would endogenously produce this pattern; however, to focus on the implications of investment for GDP growth, we assume the firm maximizes

value without frictions and then separately assume that capital adjusts gradually toward its frictionless target K^*: $K_{t+1} = K_t + (1/\theta_a)(K^* - K_t)$, with $\theta_a = 3$, beginning with the 2028 investment decision. Prior to 2028, we take investment from the calibration totals. The path starts in 2028 from the calibrated target $K^*(z_1, 0)$, so the gap that opens in 2028 is due to the rise in the boom probability alone. Panel B of Figure 1 displays the resulting paths.

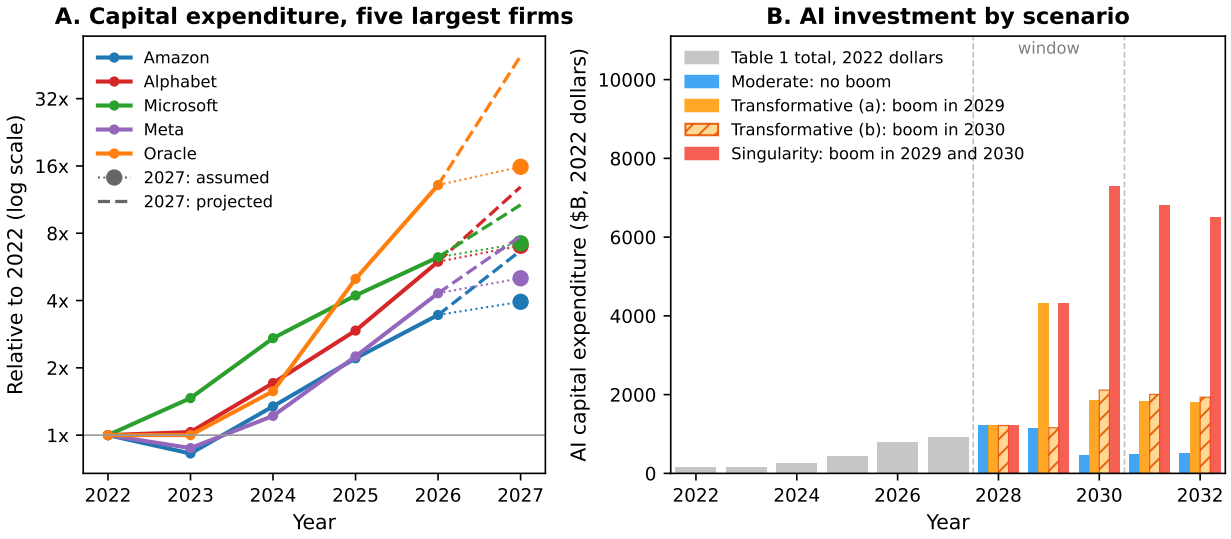


Figure 1: Capital expenditure and investment paths. Panel A: capital expenditure of the five largest firms relative to 2022 (log scale). Solid lines show Table 1, 2022–2026, in current dollars. Filled circles show the assumed 2027 value, in 2026 dollars, which keeps each firm's net addition to its capital stock, measured in 2022 dollars, at its 2026 level and is used in the calibration; dashed lines show the projected 2027 value, which continues the year-over-year growth of the firm's latest reported quarter. Panel B: AI investment by scenario, in 2022 dollars. Gray bars show AI capital expenditure through 2026 and the assumed 2027 value (the calibration totals of Table 1). Colored bars show investment implied by partial adjustment toward the frictionless target K^*: each year the capital stock closes a fraction $1/\theta_a$ of the gap, with $\theta_a = 3$, beginning with the 2028 response to the elevated boom probability. Dashed lines delineate the high-probability window (2028–2030).

Booms after 2030. What happens after 2030 is more speculative. AI is an innovation that raises the long-run growth rate, by an amount we do not know. We assume that booms of the same size continue to arrive with probability p_* per year, so that expected AI-sector productivity growth is $p_*(e^{\xi} - 1)$. The value of the firm is finite only if this growth rate

is below r_a, which requires $p_* < r_a/(e^{\xi} - 1) = 0.0351$. We set $p_* = 0.035$, just below this bound, so that expected AI-sector growth is 7.99% per year. A slightly lower p_* gives nearly the same paths: at $p_* = 0.034$, expected cumulative GDP growth by 2050 is 30% in the moderate scenario and 348% in the singularity scenario, compared with 32% and 363% at $p_* = 0.035$ (Section 4).

Table 2: Post-window outcomes by scenario

	AI share (%)	GDP gain (%)	I_a ($B/yr)	K_a^* ($B)	Capacity (GW)
Initial	2.4	—	156	624	18
Moderate	7.5	5.6	545	2,179	62
Transformative	21.1	23.7	1,788	7,153	203
Singularity	46.7	83.2	5,869	23,476	666

Notes: Each scenario row is the steady state after the window: productivity at the scenario's terminal value, and capital at the target K_a^* for that productivity and the boom probability $p_* = 0.035$, so that the option value of future booms is included, but none has occurred. The initial row is the economy before the first boom, with the boom probability at zero. "AI share" is the AI sector's share of output. "GDP gain" is the increase in the level of GDP from its initial value to the post-window steady state, holding non-AI output fixed and excluding normal-times growth. I_a is AI investment in a year without a boom, which replaces depreciation, $\delta_a K_a^*$. K_a^* is the target stock (8), and in the initial row the initial stock K_{a0}. I_a and K_a^* are in 2022 dollars. Capacity converts K_a^* to 2026 dollars with the consumer price index and then to gigawatts at a working value of $40 billion per GW.

Table 2 reports the post-window installed capital stocks K_a^* and the corresponding capacity in gigawatts.

3.4 Discussion

The revealed-preference argument identifies the productivity boom that managers believe has occurred. It does not establish that the boom has in fact occurred. Several channels could drive a wedge between the implied $\hat{\xi}$ and the true productivity gain.

First, there are various measurement issues. The 2022 baseline includes non-AI investment; if that part grew more slowly than the AI part, the initial AI capital stock is smaller

than assumed and the boom larger. Capital expenditure that the data miss, purchases on credit until they are paid (Appendix A) and firms outside Table 1, falls mainly in the later years and also makes the boom larger. The boom is also larger if firms restrain their investment to preserve monopoly power. Investment made in anticipation of the rise in p_t would make the boom smaller (Section 3.3), as would a lower depreciation rate (Appendix C).

Second, the observed investment may reflect a bubble. Managers are not immune to collective over-optimism, and the history of technology investment is replete with episodes, such as the late-1990s fiber-optic buildout, in which firms invested heavily in capacity that was not justified by subsequent demand. The productivity-enhancing power of AI may be a mirage, in which case customers would be unwilling to pay for more AI usage. In terms of the model, this is a mis-estimation of ξ, resulting in negative NPV investments, declining stock valuations, failure to make interest payments on debt, and perhaps bankruptcy. The resulting disinvestment would then lower the GDP growth estimates in Section 4.

Third, the structure of the AI infrastructure market may lead to over-investment. Dasgupta and Stiglitz (1980) show that in R&D races with winner-take-all payoffs, each firm invests as if it may capture a disproportionate share of the market, producing aggregate investment in excess of the cooperative optimum. The firms in Table 1 are building infrastructure to serve overlapping markets, and each firm's capital expenditure may reflect an expectation of capturing a larger share than is collectively feasible. If so, aggregate investment overstates the productivity gain.

Suppose, however, that each firm were to act as if it will be the sole survivor in 2027, with all investment bespoke to that firm's own AI program. The first-order condition then applies separately to each firm's own capital stock, and the revealed-preference calibration delivers a firm-specific $\hat{\xi}_i$, with $e^{\hat{\xi}_i} = K_{i,2028}/K_{i0}$. Because the aggregate capital stock is the sum of the firms', $e^{\hat{\xi}}$ is an average of the $e^{\hat{\xi}_i}$, weighted by each firm's share of the initial capital stock, so a firm-by-firm calibration implies the same productivity jump on average.

If not all firms survive, installed capital will be lower, reducing GDP growth. This prospect also cuts the other way: if firms anticipate acquiring rivals' capital at below replacement cost should those rivals fail, the option to buy later reduces the incentive to build now, and the estimate understates the boom.

Finally, just as over-optimism could lead to over-investment, managerial risk aversion, anchoring to the status quo, or the desire for the quiet life could lead to under-investment. Our revealed-preference argument would then be conservative. Indeed, our numbers are conservative compared with some industry forecasts, which align with the singularity scenario or sit above it, implying that future booms are viewed not as the high end but as the median. SpaceX's S-1 (Space Exploration Technologies Corp., 2026), filed in May 2026, sets a goal of launching 100 GW of orbital AI compute capacity per year, with first deployments "as early as 2028"; sustained, this is a steady-state stock of 400 GW for SpaceX alone, and SpaceX would not be the sector's only provider. The whole sector's stock in the singularity scenario is roughly 670 GW (Table 2). Amodei (2026) envisions "millions of instances" of AI running at "10x–100x human speed," a "country of geniuses in a datacenter," arriving in roughly 2028, which enters the model as an addition to effective labor (Appendix B). Under one interpretation, 10 million instances at 30 times human speed, this is 300 million effective workers, added to the AI sector's 13 million at the end of 2027 (2.4% of the 167 million U.S. labor force, grown by the first boom): a factor of roughly 24, against roughly 11 for the two additional booms of the singularity scenario, and so well beyond it.

Ultimately, the model's predictions should be read as conditional on the revealed-preference estimate of $\hat{\xi}$, which is itself conditional on the assumption that the firms driving the buildout are, in aggregate, not destroying value.

4 GDP growth

We now derive implications for GDP growth. In each scenario, non-AI output is held at its initial level Y_{n0}. Cumulative GDP growth from AI at time t is $(Y_{at} + Y_{n0})/\text{GDP}_0 - 1$, where $\text{GDP}_0 = Y_{a0} + Y_{n0}$ is initial GDP. A scenario's post-window steady state (Table 2) is the state in which capital has adjusted to the scenario's terminal productivity and further booms remain possible at the rate p_*, but none has occurred.

The moderate scenario's post-window steady state has a simple interpretation. AI initially accounts for about 2.4% of GDP, and its output rises by a factor of about 3.4 in the post-window steady state, $e^{\hat{\xi}} \approx 3.3$ from the boom and slightly more because the boom probability stays at p_* rather than zero, which raises the target capital stock. With non-AI output held fixed, GDP in the post-window steady state is therefore higher than its initial value by the AI share times the increase in AI output, about 5.6%. This is a permanent increase in the level of GDP from the first boom alone.

GDP during the transition. Panel A of Figure 2 shows cumulative GDP growth from AI for each scenario during the transition and its immediate aftermath (2024–2032). In the moderate scenario, GDP reaches its post-window steady state, 5.6% above its initial value, by 2030. It is temporarily about 7% above its initial value and then falls back to the steady state: the capital built during the window, when the boom probability was elevated, exceeds the target once no further boom occurs, and the excess depreciates. In the transformative scenario, with the second boom in 2029, GDP is about 23% above its initial value by 2030, essentially its post-window steady-state value. In the singularity scenario, two additional booms compound with the first to raise GDP by about 53% by 2030, against 83% in the post-window steady state. In that scenario, AI investment in a year without a boom, \$5,869B (Table 2), exceeds all U.S. private fixed investment in 2022.

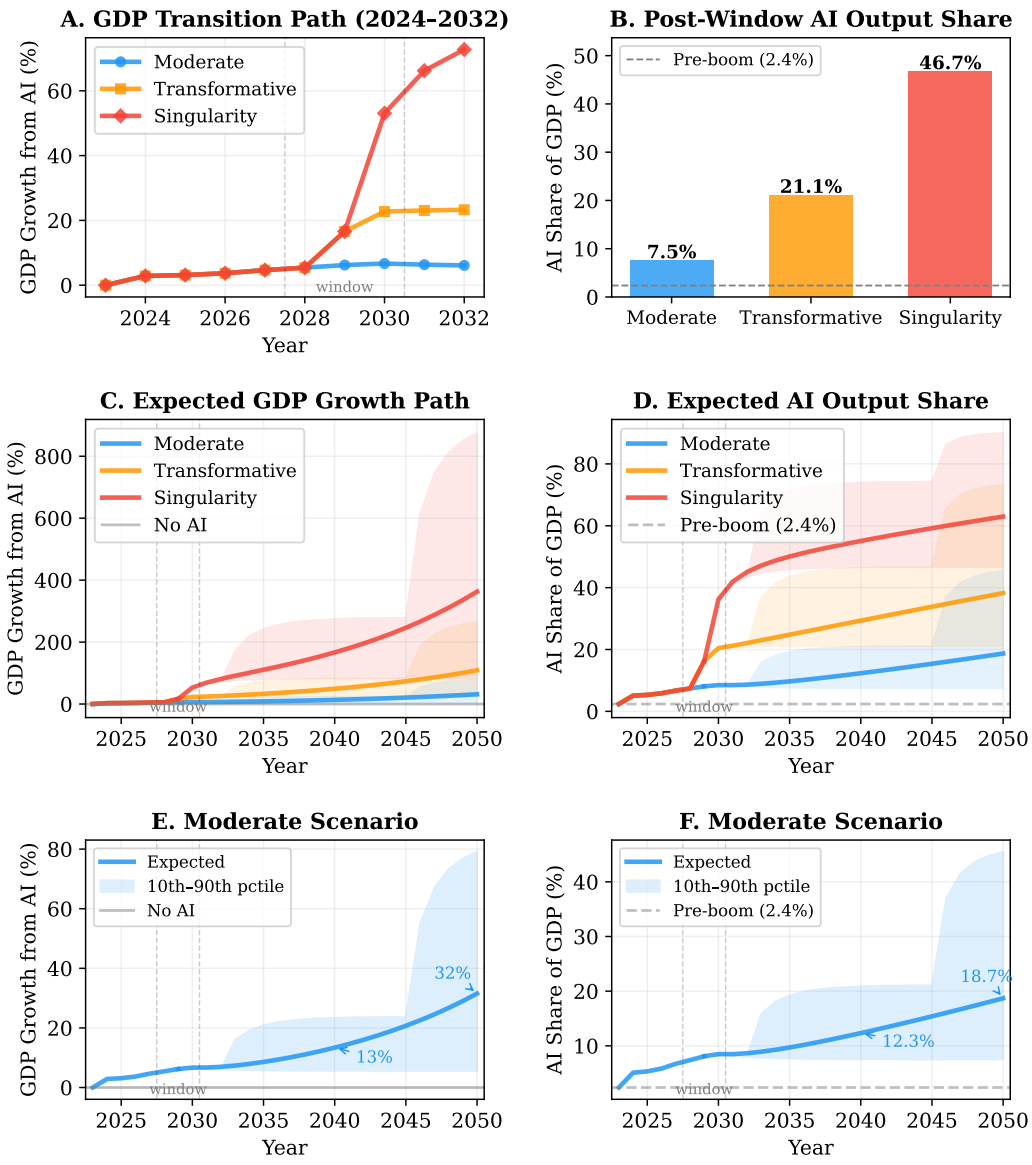


Figure 2: GDP growth and the AI output share by scenario. Panel A: cumulative GDP growth from AI during the transition and its immediate aftermath (2024–2032). Panel B: post-window AI output share of GDP. Panels C and D: expected paths to 2050 for all three scenarios. Panels E and F: the moderate scenario alone. In Panels C to F, solid lines show the expected path and shaded bands the 10th–90th percentile range, computed exactly over the 2^{20} post-window boom histories; from 2031 onward booms arrive i.i.d. at rate $p_* = 0.035$. In the GDP panels the gray line is the no-AI baseline, and in the share panels the dashed line is the initial share (2.4%). The paths use the totals in 2022 dollars of Table 1 through 2027. Beginning with the 2028 investment decision, capital follows the partial-adjustment rule of Section 3.3 ($\theta_a = 3$). The transformative scenario uses path (a), with the additional boom in 2029. Vertical dashed lines mark the 2028–2030 window.

Long-run paths (2030–2050). Panels C to F of Figure 2 extend the paths to 2050, allowing for ongoing booms. Starting from each scenario's post-window productivity level z, booms arrive i.i.d. at rate p_* from 2031, and capital continues to follow the partial-adjustment rule of Section 3.3.

In the moderate scenario, expected cumulative GDP growth reaches approximately 32% by 2050, with the AI sector's expected output share rising to roughly 19% (10th–90th percentile ranges: 5.6–79% for GDP, 7.5–45% for share). In the singularity scenario, expected cumulative GDP growth reaches 363% by 2050, with the AI sector's expected share reaching 63% of output (10th–90th percentile ranges: 83–875% for GDP, 47–90% for share).

What drives the long-run paths. Two observations summarize this discussion. First, because the model assumes no spillovers from the AI sector to the non-AI sector, the size of the AI sector after the transition is of first-order importance: aggregate growth from AI is the sector's own growth weighted by its share, and as the share approaches one, aggregate growth converges to the AI sector's growth rate. The gap between a post-window share of 7.5% (moderate) and 47% (singularity) is what compounds into the difference between 32% and 363% expected growth by 2050. Second, the number of booms after 2030 matters as much as the number in the window: the 10th–90th percentile bands of the three scenarios overlap, and the moderate scenario's 90th percentile, 79%, nearly reaches the singularity scenario's 10th, 83% (Figure 2, Panel C).

5 Historical comparisons

Several historical episodes provide precedent for the kind of productivity transformation the AI boom represents. Table 3 compares the AI scenarios with other modern sustained growth episodes, expressed as cumulative multipliers on GDP per capita.

Table 3: Historical growth episodes and the AI transition

Episode	Period	Years	Multiplier	
			GDP per capita	AI sector
Industrial revolutions				
British Ind. Rev. (1st)	1760–1840	80	1.4	—
U.S. Railroad era	1850–1910	60	2.7	—
2nd Ind. Rev. (U.S.)	1870–1920	50	2.1	—
East Asian growth miracles				
Japan postwar	1945–1970	25	4.2	—
South Korea	1960–1990	30	9.0	—
Taiwan	1960–1990	30	7.3	—
Singapore	1965–1990	25	5.7	—
China	1990–2020	30	5.8	—
Technology booms				
U.S. IT boom	1995–2005	10	1.3	—
AI window scenarios (booms through 2030, GDP in 2030)				
AI moderate (1 boom)	2024–2030	6	1.1	3.3
AI transformative (2 booms)	2024–2030	6	1.2	10.8
AI singularity (3 booms)	2024–2030	6	1.5	35.4
AI post-window, 30-year horizon (booms continue at $p_* = 0.035$)				
AI moderate	2030–2060	30	1.6	10.0
AI transformative	2030–2060	30	2.7	10.0
AI singularity	2030–2060	30	5.8	10.0

Notes: Historical multipliers are real GDP per capita in 2011 international dollars from the Maddison Project Database 2023 (Bolt and van Zanden, 2025); over the periods shown, the database's series are those of Broadberry et al. (2015) for Britain and Cha et al. (2022) for South Korea. For the AI rows, the GDP column reports the multiplier on GDP from AI alone, with non-AI output held at its initial level and normal-times growth excluded. For the window scenarios, Period and Years identify when the productivity booms occur; the GDP multipliers are the 2030 values on the partial-adjustment paths of Figure 2. In the post-window steady state they are 1.1, 1.2, and 1.8 (Table 2). The 30-year rows are expected GDP in 2060 relative to GDP in 2030 on the same paths, carried to 2060 by simulation, with booms arriving at rate p_* from 2031. A proportional increase in GDP is the same per capita. The AI-sector column reports the sector's own productivity multiplier ($e^{n\xi}$ for n window booms; post-window, its expected value $(1+g)^{30}$, with $g = p_*(e^{\xi} - 1) \approx 8.0\%$); it has no counterpart in the historical episodes.

The AI multipliers measure the contribution of AI booms alone; the historical multipliers include growth from every source, so the comparison understates the AI scenarios. By 2030, the moderate scenario raises GDP by a factor of 1.1 and the singularity scenario by 1.5, compared with 1.3 for the U.S. IT boom over a decade. On a 30-year horizon, with booms continuing at $p_* = 0.035$, the expected multipliers are 1.6 (moderate), 2.7 (transformative), and 5.8 (singularity), compared with 1.4–2.7 for the industrial revolutions over 50–80 years and 4.2–9.0 for the East Asian miracles over 25–30 years.

The U.S. railroad era provides perhaps the closest historical analogue to the current AI buildout. Route mileage rose from approximately 9,000 miles in 1850 to 164,000 in 1890 and continued growing through 1916, peaking at 254,000 miles.[7] By 2022, the national railroad network had contracted to approximately 136,000 route miles, about 46% below its 1916 peak (Bureau of Transportation Statistics, 2024, p. 1-24).

This contraction unfolded over many decades and does not by itself establish that the original infrastructure investment was wasteful. GDP per capita rose 2.7× over the railroad era (Table 3). The original productivity gain was real; the subsequent decline was a negative shock to the sector's z from a competing technology, not a revelation that the gain itself was illusory. For the AI transition, the railroad precedent suggests that transformative booms can involve capital-stock overshooting, but only after long intervals and without negating the underlying productivity gain. The relevant question is whether the productivity boom justifies the investment at the time it is made—not whether some fraction of today's GPU capacity will eventually be stranded, since depreciation at $\delta_a = 0.25$ already implies that current hardware will be largely replaced within four years.

Railroad capital expenditures reached roughly 5–6% of GNP in the peak years of the early 1870s and absorbed about one-fifth of gross domestic fixed capital formation over

[7] Historical Statistics of the United States, Series Df927; Association of American Railroads.

1870–1874.[8] Private fixed investment in communication equipment and structures peaked at 1.2% of GDP in 1999–2000, the largest communications share in the BEA series.[9] The AI shares in Panel B of Table 1 are 2.8% of GDP in 2026, more than twice the telecom peak, and, at 15.5% of private fixed investment, approach the railroads' one-fifth of capital formation.

6 Conclusion

In our two-sector model, only the AI sector is exposed to rare productivity booms. The capital stock built by investment through 2026, with building continuing at its 2026 pace in 2027, implies that the boom raised AI-sector productivity by a factor of roughly 3.3. This boom alone (the moderate scenario) implies a permanent increase in the level of GDP of approximately 5.6%. In the transformative and singularity scenarios, one and two further booms arrive during a two-year window of elevated boom probability, and by 2030 GDP is approximately 23% and 53% above its initial level, respectively. These GDP gains hold non-AI output fixed and exclude normal-times growth. The AI sector's share of the economy rises from 2.4% before the boom to between 7.5% and 47% depending on the scenario, once capital has fully adjusted.

AI raises the long-run growth rate, by an amount the investment data do not determine. If booms continue at the largest annual probability at which firm values are finite, expected annual growth of the AI sector is approximately 8%.

There is a vast range of scenarios, and then substantial long-run uncertainty built on top

[8] Computed from railroad gross capital expenditures in Ulmer (1960), Table C-1, and the nominal GNP series of Gallman and Rhode (2020). Because the annual interpolation underlying single-year values is imprecise, the five-year ratio is the more reliable number: cumulative 1870–1874 capital expenditure relative to cumulative GNP is approximately 4%. Fishlow (1966) argues that replacement accounting leads Ulmer's series to understate gross investment, so these shares are, if anything, lower bounds.

[9] Bureau of Economic Analysis: private fixed investment in communication equipment ($103B) plus communication structures ($19B) in 2000, relative to GDP of $10.25T; the 1999 share is nearly identical.

of those. At its base, however, is a productivity boom that is in investment data, but not yet in productivity data. If the boom fails to materialize, the current buildout will be the largest misallocation of capital in history. On the other hand, those who are certain of that outcome may be making a mistake analogous to that which many forecasters made in 2007, but in reverse. Rare events are hard to imagine until they occur.

References

Acemoglu, D. (2024). The simple macroeconomics of AI. *NBER Working Paper 32487*.

Alphabet Inc. (2023). Alphabet announces first quarter 2023 results. Earnings release, April 25, 2023; Exhibit 99.1 to Form 8-K. Available at https://www.sec.gov/Archives/edgar/data/1652044/000165204423000041/googexhibit991q12023.htm.

Amazon.com, Inc. (2023). 2022 annual report on form 10-K. U.S. Securities and Exchange Commission, fiscal year ended December 31, 2022. Available at https://www.sec.gov/Archives/edgar/data/1018724/000101872423000004/amzn-20221231.htm.

Amazon.com, Inc. (2025). 2024 annual report on form 10-K. U.S. Securities and Exchange Commission, fiscal year ended December 31, 2024. Available at https://www.sec.gov/Archives/edgar/data/1018724/000101872425000004/amzn-20241231.htm.

Amodei, D. (2024). Machines of loving grace: How AI could transform the world for the better. Essay, October 2024. Available at https://darioamodei.com/essay/machines-of-loving-grace.

Amodei, D. (2026). We are near the end of the exponential. Interview by Dwarkesh Patel, *Dwarkesh Podcast*, February 13, 2026. Transcript available at https://www.dwarkesh.com/p/dario-amodei-2.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866.

Bolt, J. and van Zanden, J. L. (2025). Maddison-style estimates of the evolution of the world economy: A new 2023 update. *Journal of Economic Surveys*, 39(2):631–671.

Broadberry, S., Campbell, B. M. S., Klein, A., Overton, M., and van Leeuwen, B. (2015). *British Economic Growth, 1270–1870*. Cambridge University Press.

Bureau of Transportation Statistics (2024). Transportation statistics annual report 2024. Statistical report, U.S. Department of Transportation, Washington, DC. Available at https://doi.org/10.21949/e0kq-gf72.

Cha, M. S., Kim, N. N., Park, K.-J., and Park, Y., editors (2022). *Historical Statistics of Korea*. Studies in Economic History. Springer, Singapore.

Chen, A. Y. (2026). Hedging the singularity. *Working Paper, Federal Reserve Board*. Available at https://arxiv.org/abs/2604.16997.

Cooley, T. F. and Prescott, E. C. (1995). Economic growth and business cycles. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*, pages 1–38. Princeton University Press.

Dasgupta, P. and Stiglitz, J. (1980). Industrial structure and the nature of innovative activity. *Economic Journal*, 90(358):266–293.

Eberly, J. and Wang, N. (2025). Reallocating and pricing illiquid capital: Two productive trees. *International Economic Review*, 66(5):1887–1907.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

Fishlow, A. (1966). Productivity and technological change in the railroad sector, 1840–1910. In Brady, D. S., editor, *Output, Employment, and Productivity in the United States after 1800*, pages 583–646. Columbia University Press for the NBER, New York.

Gallman, R. E. and Rhode, P. W. (2020). *Capital in the Nineteenth Century*. University of Chicago Press, Chicago.

Gomes, J. F., Grotteria, M., and Wachter, J. A. (2019). Cyclical dispersion in expected defaults. *Review of Financial Studies*, 32(4):1275–1308.

Gourio, F. (2012). Disaster risk and business cycles. *American Economic Review*, 102(6):2734–2766.

Hanson, R. (2000). Long-term growth as a sequence of exponential modes. Working paper, George Mason University.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):213–224.

Kording, K. and Marinescu, I. (2025). (Artificial) intelligence saturation and the future of work. Working paper, Brookings Institution, Center on Regulation and Markets. Available at https://www.brookings.edu/articles/artificial-intelligence-saturation-and-the-future-of-work/.

Korinek, A. and Suh, D. (2024). Scenarios for the transition to AGI. *NBER Working Paper 32255*.

Meta Platforms, Inc. (2025). 2024 annual report on form 10-K. U.S. Securities and Exchange Commission, fiscal year ended December 31, 2024. Available at https://www.sec.gov/Archives/edgar/data/1326801/000132680125000017/meta-20241231.htm.

Microsoft Corporation (2023). Annual report 2023. Annual report for the fiscal year ended June 30, 2023. Available at https://www.microsoft.com/investor/reports/ar23/index.html.

Miller, M., Paron, J. D., and Wachter, J. A. (2025). Sovereign default and the decline in interest rates. *Review of Financial Studies*. Forthcoming.

Restrepo, P. (2025). We won't be missed: Work and growth in the AGI world. *NBER Working Paper 34423*.

Space Exploration Technologies Corp. (2026). Form S-1 registration statement. U.S. Securities and Exchange Commission, filed May 20, 2026, accession 0001628280-26-036936. Available at https://www.sec.gov/Archives/edgar/data/1181412/000162828026036936/spaceexplorationtechnologi.htm.

Trammell, P. and Korinek, A. (2023). Economic growth under transformative AI. *NBER Working Paper 31815*.

Tsai, J. and Wachter, J. A. (2016). Rare booms and disasters in a multisector endowment economy. *Review of Financial Studies*, 29(5):1113–1169.

Ulmer, M. J. (1960). *Capital in Transportation, Communications, and Public Utilities: Its Formation and Financing*. Princeton University Press, Princeton.

Van Nieuwerburgh, S. (2026). Financing the AI buildout. *Journal of Economic Perspectives*. Forthcoming; draft of March 20, 2026.

Vanguard Group (2025). AI exuberance: Economic upside, stock market downside. Vanguard Research. Available at https://corporate.vanguard.com/content/corporatesite/us/en/corp/vemo/ai-exuberance-economic-upside-stock-market-downside.html.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68(3):987–1035.

Yotzov, I., Barrero, J. M., Bloom, N., et al. (2026). Firm data on AI. *NBER Working Paper 34836*.

A Capital expenditure data

A.1 Capital expenditure of the five largest firms, 2022–2025

Calendar-year basis. Amazon, Alphabet, and Meta report on the calendar year (Table A.1). Microsoft's fiscal year ends June 30 and Oracle's ends May 31. Microsoft's fiscal quarters are calendar quarters, and it states its capital expenditure for each quarter on its earnings call, so its figure for each calendar year is the sum of four quarters. For Oracle, let fiscal year T be the fiscal year that ends in calendar year T. Capital expenditure for year T is capital expenditure in

$$\text{fiscal year } T \;-\; \text{first six months of fiscal year } T \;+\; \text{first six months of fiscal year } T+1.$$

Oracle's fiscal quarters end in August, November, February, and May, so its filings contain no figure for a period that ends in December. The first six months of Oracle's fiscal year are June through November, so the construction gives December of year $T-1$ through November of year T. The same construction applied to Microsoft's statements of cash flows, whose first six months run from July through December, gives January through December of year T. Table A.2 reports the figures and the construction, and Figure A.1 shows the construction for 2025.

Amazon. Amazon's measure of its capital expenditure, in the free-cash-flow reconciliation of its 10-K and its earnings releases, is "Purchases of property and equipment, net of proceeds from sales and incentives": the cash-flow line "Purchases of property and equipment" less the line "Proceeds from property and equipment sales and incentives" (Table A.1). Table 1 counts all of Amazon's capital expenditure, including spending outside AWS. Amazon reports net additions to property and equipment by segment; AWS accounted for 51% of them in

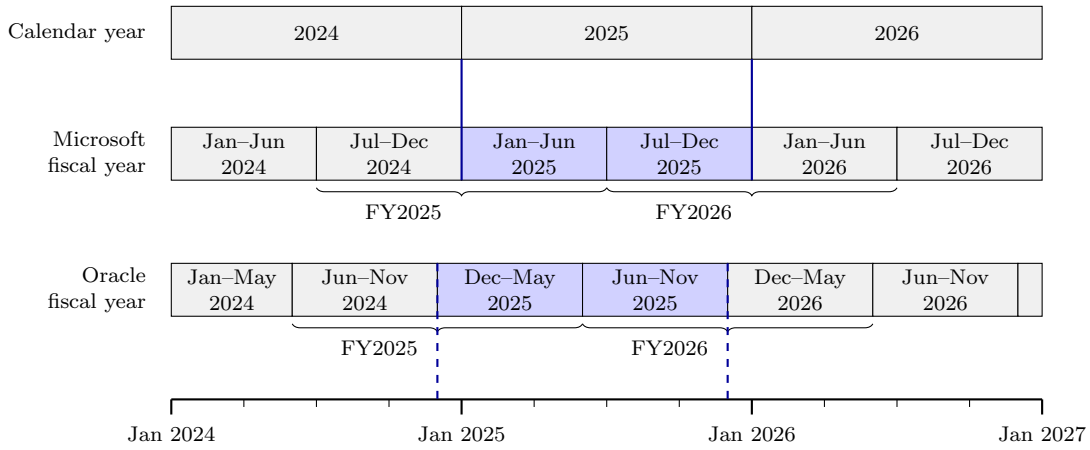


Figure A.1: Fiscal years and the calendar-year construction for Microsoft and Oracle. The 2025 figure in Table 1 is the sum of the shaded half-years: the second half of fiscal 2025 and the first half of fiscal 2026. For Microsoft they cover January–December 2025 (between the solid lines). For Oracle they cover December 2024–November 2025 (between the dashed lines), one month earlier than the calendar year.

Table A.1: Capital expenditure of Amazon, Alphabet, and Meta ($ millions, current dollars)

	2022	2023	2024	2025
Amazon: "Purchases of property and equipment"	63,645	52,729	82,999	131,819
less: proceeds from asset sales and incentives	5,324	4,596	5,341	3,499
equals: net purchases, as Amazon measures them†	58,321	48,133	77,658	128,320
Alphabet: "Purchases of property and equipment"†	31,485	32,251	52,535	91,447
Meta: "Purchases of property and equipment"	31,186	27,045	37,256	69,691
plus: principal payments on finance leases	850	1,058	1,969	2,524
equals: capital expenditures as Meta defines them†	32,036	28,103	39,225	72,215

Notes: Figures are from the statements of cash flows in each firm's Form 10-K, except the rows labeled *equals*, which are computed. † The 2022–2025 entries in Table 1 are this row, rounded to the nearest billion dollars.

Table A.2: Capital expenditure of Microsoft and Oracle by calendar year (current dollars)

Year T	2022	2023	2024	2025	2026
Microsoft: capital expenditures including finance leases, from the earnings calls ($ billions)					
January–March	6.3	7.8	14.0	21.4	31.9
April–June	8.7	10.7	19.0	24.2	41.0
July–September	6.6	11.2	20.0	34.9	—
October–December	6.8	11.5	22.6	37.5	—
January–December[†]	28.4	41.2	75.6	118.0	—
Memo: additions to property and equipment[‡]	24.8	35.2	55.6	83.1	—
Oracle: "Capital expenditures" ($ millions); fiscal year T is June $T-1$–May T					
(a) Fiscal year T	4,511	8,695	6,866	21,215	55,663
(b) June–November $T-1$	1,987	4,154	2,394	6,273	20,535
(c) June–November T	4,154	2,394	6,273	20,535	—
(a) − (b): December $T-1$–May T	2,524	4,541	4,472	14,942	35,128
(a) − (b) + (c): December $T-1$–November T[†]	6,678	6,935	10,745	35,477	—

Notes: Microsoft: each quarter's figure is the one its chief financial officer states on the earnings call for that quarter; Microsoft's fiscal quarters are calendar quarters. [‡]Additions to property and equipment, January–December, from the statements of cash flows, constructed as for Oracle from the Forms 10-K and the Forms 10-Q for the second fiscal quarter; $66.7B in January–June 2026. Oracle: rows (a)–(c) are from the statements of cash flows. Row (a) is from the Form 10-K. Rows (b) and (c), the first six months of fiscal years T and $T+1$, are from the Forms 10-Q for the second quarters of those years. A dash indicates that the call or the Form 10-Q for the quarter had not taken place or been filed by September 29, 2026. [†] The 2022–2025 entries in Table 1 are this row, rounded to the nearest billion dollars; Appendix A.2 describes the 2026 entries.

2023, 62% in 2024, and 68% in 2025 (10-K for 2025, segment information). Amazon's guidance for 2026 is for the whole company. The measure excludes equipment acquired under finance leases, which Amazon reports as a non-cash supplemental disclosure, and counts equipment bought on credit when it is paid for: property and equipment acquired but not yet paid fell by $1.4B in 2023 and rose by $7.0B in 2024 and $10.2B in 2025 (10-K for 2025, supplemental cash-flow information). Sources: 2022 figures from the 10-K for 2024 (accession 0001018724-25-000004, filed February 7, 2025); 2023–2025 figures from the 10-K for 2025 (accession 0001018724-26-000004, filed February 6, 2026).

Alphabet. The cash-flow line is "Purchases of property and equipment," and Alphabet's 10-K reports its capital expenditures as this line ($91.4B in 2025). Alphabet does not include finance-lease additions in this line and does not subtract proceeds from sales of property and equipment. Sources: 2022 figures from the 10-K for 2024 (accession 0001652044-25-000014, filed February 5, 2025); 2023–2025 figures from the 10-K for 2025 (accession 0001652044-26-000018, filed February 5, 2026).

Microsoft. Microsoft's chief financial officer states capital expenditures including finance leases for each quarter on the earnings call. Equipment counts when it is received rather than when it is paid for, and finance leases count at their full value when the lease begins ("where we recognize the full value at the time of lease commencement," call of July 30, 2025). On the call of April 29, 2026, capital expenditures were $31.9B for January–March 2026, finance leases $4.7B, and cash paid for property and equipment $30.9B; the finance leases were partly offset by "differences between the receipt of goods and payment." Microsoft does not report the figure in its earnings releases, its 10-Qs, or its 10-K. Sources: transcripts of the earnings calls for the third quarter of fiscal 2022 through the fourth quarter of fiscal 2026 (Microsoft Investor Relations). The cash-flow line in the memo row of Table A.2, "Additions

to property and equipment," excludes assets acquired under finance leases ($24.6B in fiscal 2026), which Microsoft reports as a non-cash supplemental disclosure, and principal payments on finance leases, which it reports as financing cash flows (10-K for fiscal 2026, Note 13). Its sources are: fiscal 2022 figures from the 10-K for fiscal 2024 (accession 0000950170-24-087843, filed July 30, 2024); fiscal 2023–2025 figures from the 10-K for fiscal 2025 (accession 0000950170-25-100235, filed July 30, 2025); fiscal 2026 figures from the 10-K for fiscal 2026 (accession 0001193125-26-323660, filed July 29, 2026). Six-month figures are from the 10-Q for the quarter ended December 31 of each year: 2021 (accession 0001564590-22-002324, filed January 25, 2022); 2022 (accession 0001564590-23-000733, filed January 24, 2023); 2023 (accession 0000950170-24-008814, filed January 30, 2024); 2024 (accession 0000950170-25-010491, filed January 29, 2025); 2025 (accession 0001193125-26-027207, filed January 28, 2026).

Meta. Meta defines its capital expenditures as the sum of two cash-flow lines: "Purchases of property and equipment" and "Principal payments on finance leases," a financing cash flow. Table 1 uses this sum in every year. Sources: 2022 figures from the 10-K for 2024 (accession 0001326801-25-000017, filed January 30, 2025); 2023–2025 figures from the 10-K for 2025 (accession 0001628280-26-003942, filed January 29, 2026).

Oracle. The cash-flow line is "Capital expenditures." Assets acquired under finance leases are not in this line. The line includes equipment that Oracle buys with customer prepayments and excludes GPUs that customers buy and supply to Oracle under its large AI contracts (earnings release, Form 8-K, March 10, 2026, Exhibit 99.1). Oracle does not report the customer-supplied amount, so the line understates the equipment installed in Oracle's data centers. The 10-Q for each quarter ended November 30 reports capital expenditures for the trailing four quarters in the free-cash-flow table of management's discussion and analy-

sis. These equal the December–November figures constructed in Table A.2. Sources: fiscal 2022 figures from the 10-K for fiscal 2024 (accession 0000950170-24-075605, filed June 20, 2024); fiscal 2023–2025 figures from the 10-K for fiscal 2025 (accession 0000950170-25-087926, filed June 18, 2025); fiscal 2026 figures from the 10-K for fiscal 2026 (accession 0001193125-26-277521, filed June 22, 2026). Six-month figures are from the 10-Q for the quarter ended November 30 of each year: 2021 (accession 0001564590-21-060022, filed December 10, 2021); 2022 (accession 0001564590-22-039278, filed December 13, 2022); 2023 (accession 0000950170-23-069682, filed December 12, 2023); 2024 (accession 0000950170-24-134973, filed December 10, 2024); 2025 (accession 0001193125-25-315925, filed December 11, 2025).

Leases. Table A.3 reports the assets the five firms acquired under leases, which their capital expenditure leaves out except for Microsoft's finance leases and Meta's principal payments. In the latest year, operating-lease additions were 4 to 15% of cash purchases of property and equipment at Amazon, Alphabet, Meta, and Microsoft, down from 9 to 27% two or three years earlier, and a third at Oracle, which leases much of its data-center capacity. We keep each firm's own measure so that its forward guidance, which is on that measure, can be used. Counting leased assets would raise every year's figure, later years by more dollars and earlier years by a larger share, so the effect on the boom size is ambiguous and, at these shares, small. Equipment financed by suppliers or lenders, as at CoreWeave and xAI, is likewise paid for through financing cash flows and enters the cash measures only as it is paid.

Table A.3: Assets acquired under leases, relative to cash purchases of property and equipment ($ billions, current dollars)

		Purchases of	Operating leases		Finance leases	
		property and equipment	$B	%	$B	%
Amazon	2023	52.7	14.1	27	0.6	1
	2024	83.0	15.4	19	0.9	1
	2025	131.8	19.9	15	2.9	2
Alphabet	2023	32.3	2.9	9	0.6	2
	2024	52.5	2.5	5	0.3	1
	2025	91.4	4.1	4	1.6	2
Meta	2022	31.2	4.4	14	0.2	1
	2023	27.0	4.4	16	0.6	2
	2024	37.3	3.8	10	0.2	0
	2025	69.7	7.0	10	0.6	1
Microsoft	FY2022	23.9	5.3	22	4.2	18
	FY2023	28.1	3.5	13	3.1	11
	FY2024	44.5	6.7	15	11.6	26
	FY2025	64.6	7.8	12	20.5	32
	FY2026	115.9	4.6	4	24.6	21
Oracle	FY2023	8.7	2.0	23		
	FY2024	6.9	4.2	62	0.0	0
	FY2025	21.2	7.0	33	2.9	14
	FY2026	55.7	18.2	33	4.9	9

Notes: Right-of-use assets obtained in exchange for lease obligations, from the leases notes and supplemental cash-flow information of the Forms 10-K (Amazon: assets acquired under operating leases, and property and equipment acquired under finance leases, net of remeasurements and modifications), as a percentage of the cash-flow line for purchases of property and equipment (Microsoft: additions to property and equipment; Oracle: capital expenditures). Fiscal years end June 30 for Microsoft and May 31 for Oracle. Microsoft's capital expenditures include its finance leases; Meta's include the principal payments on its finance leases (Table A.1). Oracle's finance-lease additions in fiscal 2023 are not disclosed. Sources: Forms 10-K of Amazon, Alphabet, and Meta for 2025, Meta for 2024, Microsoft for fiscal 2022, 2025, and 2026, and Oracle for fiscal 2025 and 2026.

A.2 Capital expenditure in 2026 and the 2027 values used in the calibration

Table A.4 reports the 2026 figures of the five largest firms, from filings available on EDGAR through September 29, 2026, and from Microsoft's earnings calls.

Table A.4: Capital expenditure of the five largest firms in 2026 ($ millions, current dollars)

	Actual	Full year			Table 1
	to date	Low	Point	High	($ billions)
Amazon	96,310		200,000		200
Alphabet	80,598	180,000	185,000	190,000	185
Microsoft	72,900		175,000		175
Meta	50,918	130,000	137,500	145,000	138
Oracle	63,627	86,123	92,126	102,189	92
Total, five largest firms			789,626		790

Notes: "Actual to date" is capital expenditure in January–June 2026 (Oracle: December 2025–August 2026), measured as for 2022–2025 (Appendix A.1). "Full year" is January–December 2026 (Oracle: December 2025–November 2026). The full-year figures are management's, except Oracle's, which are the actual to date plus an extrapolation for the remaining months. For Alphabet and Meta, the point figure is the midpoint of management's range. Microsoft's actual to date is the sum of the two quarters' figures stated on its calls. The last column is the point figure rounded to the nearest billion dollars; its total is the sum of the rounded entries.

2026 figures

Amazon. The earnings release for the fourth quarter of 2025 (Form 8-K, February 5, 2026, Exhibit 99.1; accession 0001018724-26-000002) states that Amazon expects "to invest about $200 billion in capital expenditures across Amazon in 2026." No later filing changes this figure. On Amazon's measure, used for 2022–2025, capital expenditure was $96.310B in January–June 2026: $98.411B of purchases of property and equipment less $2.101B of proceeds from property and equipment sales and incentives (10-Q filed July 31, 2026; accession 0001018724-26-000026). The February 5 release does not say whether the $200 billion is

before or after these proceeds; we take it to be on Amazon's measure.

Alphabet. The prospectus supplement for Alphabet's offering of Class A and Class C stock (Form 424B5, filed June 4, 2026; accession 0001193125-26-256375) states under "Use of Proceeds": "We expect that capital expenditures in 2026 will be in the range of $180 billion to $190 billion, and that 2027 capital expenditures will significantly increase compared to 2026." The press release furnished with the Form 8-K of June 4, 2026 (Exhibit 99.1; accession 0001193125-26-257724) states that Alphabet announced both expectations on its earnings call for the first quarter of 2026, held April 29, 2026. Alphabet's earnings releases for the first and second quarters of 2026 give no figure. Purchases of property and equipment were $80.598B in January–June 2026 and $39.643B in January–June 2025 (10-Q filed July 23, 2026; accession 0001652044-26-000071).

Microsoft. Microsoft does not issue annual capital-expenditure guidance in its earnings releases, its 10-Qs, or its 10-K; it gives its outlook on its earnings calls. On the call of April 29, 2026, its chief financial officer said that Microsoft expected to invest roughly $190B in capital expenditures in calendar 2026. On the call of July 29, 2026, she said that from the start of fiscal 2027 Microsoft is extending the estimated useful lives of its data centers from 15 to 25 years, so that more of its future data-center leases will be operating leases. Its capital expenditures include finance leases and exclude operating leases. Apart from this change, its expectation for calendar 2026 was unchanged, and the change "adjusts our expectation to approximately $175 billion" (transcripts of both calls, Microsoft Investor Relations). In January–June 2026, capital expenditures including finance leases, the measure used for 2022–2025, were $72.9B ($31.9B and $41.0B in the two quarters, as stated on the calls), and additions to property and equipment were $66.678B (Table A.2). On the July call, Microsoft gave an outlook of over $50B for July–September 2026; a presentation furnished

with the Form 8-K of September 2, 2026 restates it (Exhibit 99.1; accession 0001193125-26-380280). Microsoft's latest quarter in Table A.5 is capital expenditures including finance leases of $41.0B in April–June 2026, against $24.2B a year earlier (calls of July 29, 2026 and July 30, 2025); additions to property and equipment were $35.802B (earnings release, Form 8-K, July 29, 2026, Exhibit 99.1; accession 0001193125-26-323632).

Meta. The earnings release for the second quarter of 2026 (Form 8-K, July 29, 2026, Exhibit 99.1; accession 0001628280-26-050596) states: "We anticipate 2026 capital expenditures, including principal payments on finance leases, to be in the range of $130–145 billion." Capital expenditures were $50.918B in January–June 2026 ($49.113B of purchases of property and equipment plus $1.805B of finance-lease principal) and $30.704B in January–June 2025 (10-Q filed July 30, 2026; accession 0001628280-26-050705).

Oracle. Capital expenditures were $35.128B in December 2025–May 2026 (Table A.2) and $28.499B in June–August 2026 (10-Q for the quarter ended August 31, 2026, filed September 11, 2026; accession 0001193125-26-389274), so the total for December 2025–August 2026 is $63.627B.

Oracle has not stated a dollar figure for its fiscal 2027 capital expenditures in a filed document. For September–November 2026, we set capital expenditures equal to those of June–August 2026, $28.499B, so the point figure is $63.627B + $28.499B = $92.126B.

The low and high figures in Table A.4 also use capital expenditures of $8.502B in June–August 2025 and $16.493B in March–May 2026 (earnings release for the first quarter of fiscal 2027, Form 8-K, September 10, 2026, Exhibit 99.1; accession 0001193125-26-387905). The low figure sets September–November 2026 equal to the average of the two most recent quarters, ($16.493B+$28.499B)/2 = $22.496B, which gives $86.1B for the full year. The high figure grows capital expenditures for one more quarter at the compound quarterly growth

rate from June–August 2025 to June–August 2026, $(28.499/8.502)^{1/4} - 1 = 35\%$, which gives \$38.6B for September–November 2026 and \$102.2B for the full year.

2027 values used in the calibration

No firm has stated a dollar figure for its 2027 capital expenditure in a filed document. The calibration uses, for each firm, the 2027 capital expenditure that keeps its net addition to the capital stock, measured in 2022 dollars, at its 2026 level: 2026 capital expenditure plus depreciation on the net addition in 2026. Table A.5 reports this calculation, in 2022 dollars, and the projection in the last column of Table 1.

Table A.6 reports what management has said about capital expenditure and capacity. Table A.7 reports leases signed but not yet commenced, mostly for data centers. The leases enter capital expenditure only through Meta's principal payments on finance leases and Microsoft's finance leases (Appendix A.1); the servers and network equipment that the firms buy for the leased space enter it in full. Table A.8 reports other multi-year commitments.

Table A.6: Statements on capital expenditure and capacity

		\$ billions, current dollars		
Firm	Date	Latest	Year earlier	Statement (quantity in parentheses)
Amazon	Feb. 5, 2026	about 200	128.3	"... we expect to invest about \$200 billion in capital expenditures across Amazon in 2026" (capital expenditures in 2026; year earlier: cash capital expenditures in 2025)
	Jul. 31, 2026	96.3	55.6	"... investments in technology infrastructure (the majority of which is to support AWS business growth) and in additional capacity to support our fulfillment network, both of which we expect to increase in 2026." (cash capital expenditures, January–June)
Alphabet	Feb. 5, 2026			"Our ability to scale our technical infrastructure is increasingly constrained by the availability of power, water, and land. ... Failure to secure sufficient capacity in a timely manner would limit our ability to train models and serve Cloud customers."

Table A.6, continued

Firm	Date	$ billions, current dollars: Latest	$ billions, current dollars: Year earlier	Statement (quantity in parentheses)
	Jun. 1, 2026			"...strong demand for its AI solutions and services from enterprises and consumers, at levels that are exceeding the company's available supply."
	Jun. 2, 2026	180–190	91–93	"...customer demand for AI compute capacity, which is outstripping our current supply. We expect that capital expenditures in 2026 will be in the range of $180 billion to $190 billion, and that 2027 capital expenditures will significantly increase compared to 2026." (capital expenditures in 2026; year earlier: 2025 range stated October 29, 2025)
	Jul. 23, 2026	80.6	39.6	"In 2026, we expect to significantly increase, relative to 2025, our investment in our technical infrastructure, including servers and network equipment and data centers." (capital expenditures, January–June)
Microsoft	Jul. 29, 2026	115.9	64.6	"We will continue to invest in capital expenditures to support growth in our cloud offerings and our investments in AI training and other infrastructure." (additions to property and equipment, fiscal 2026)
	Jul. 29, 2026	329.1	92.7	"...demand exceeding available capacity limits our ability to meet customer needs in a timely manner." "There are limited suppliers for certain critical device and datacenter components, and those items are in short supply." (leases, primarily for datacenters, not yet commenced, June 30)
	Sep. 2, 2026	over 50	34.9	Capital expenditures for July–September 2026, as provided on July 29, 2026: "Expected to be over $50 billion including the impact of the useful life update" (year earlier: capital expenditures including finance leases, July–September 2025)
Meta	Apr. 29, 2026	125–145		"...This reflects our expectations for higher component pricing this year and, to a lesser extent, additional data center costs to support future year capacity." (capital expenditures in 2026, raised from the prior range)

Table A.6, continued

Firm	Date	$ billions, current dollars: Latest	$ billions, current dollars: Year earlier	Statement (quantity in parentheses)
	Jul. 29, 2026	130–145	70–72	"We anticipate 2026 capital expenditures, including principal payments on finance leases, to be in the range of $130-145 billion, narrowed from our prior outlook of $125-145 billion." (year earlier: 2025 range stated October 29, 2025)
	Jul. 30, 2026	50.9	30.7	"... we have significantly increased our infrastructure investments in connection with our AI initiatives, and expect our investments to continue to increase." (capital expenditures including principal payments on finance leases, January–June)
Oracle	Jun. 22, 2026	55.7	21.2	"We expect this upward trend to continue during fiscal 2027 and in the following fiscal years ..." (capital expenditures, fiscal 2026)
	Jun. 22, 2026			"... industry supply capacity for AI accelerators, including graphics processing units, as well as memory devices, is competitive, and we at times have to accept less favorable terms with suppliers to minimize supply constraints. While this permits us to secure cloud infrastructure capacity ..."
	Sep. 10, 2026	664	455	"Customer demand for AI Cloud Training and Inferencing Services continues to grow faster than supply." (remaining performance obligations, August 31)
	Sep. 11, 2026	28.5	8.5	"We expect our capital expenditures in fiscal 2027 to be higher than fiscal 2026 as we increase our existing data center capacity and establish data centers in new geographic locations in order to meet current and expected customer demand for our cloud offerings." (capital expenditures, June–August)
CoreWeave	May 7, 2026	99.4		"This was the strongest bookings quarter in CoreWeave's history, with revenue backlog reaching nearly $100 billion." (revenue backlog, March 31)
	May 8, 2026			"We require substantial and growing capital expenditures to support our growth and respond to business challenges."
IREN	May 7, 2026			"The world is structurally short compute, and the bottleneck is delivered data center and GPU capacity"

Table A.6, continued

Firm	Date	$ billions, current dollars: Latest	$ billions, current dollars: Year earlier	Statement (quantity in parentheses)
	May 8, 2026	11.9	0.4	"The market for GPUs and related equipment remains highly competitive and supply constrained . . . " (total commitments, March 31, 2026; year earlier: June 30, 2025)
SpaceX	May 20, 2026	15.8	0.7	"During this investment period, our capital expenditures will scale as quickly as we are able to deploy power and compute . . . " (AI-segment capital expenditures, April–June)
	May 20, 2026			"Our ability to scale our data center infrastructure, which supports our AI segment, is increasingly constrained by the availability of power at economically feasible prices, long lead times, availability of materials, and changing regulatory requirements."

Notes: Statements are quoted from the documents listed below. The date is the release date for earnings releases, press releases and presentations, the date of the prospectus supplement, and the filing date for annual and quarterly reports and the registration statement. "Latest" is the quantity described in parentheses, from the same document unless stated. "Year earlier" is the same quantity twelve months earlier unless the parenthesis gives another date. A blank cell means that the statement carries no quantity or that the filings give no comparable figure. Ranges and the words "about" and "over" are management's. Amounts converted from millions are rounded to one decimal. *Amazon:* earnings release of February 5, 2026 (Form 8-K, Exhibit 99.1); 10-Q for the quarter ended June 30, 2026 (filed July 31, 2026), liquidity section. The 2025 figure is from the 2025 10-K. Amazon's cash capital expenditures are purchases of property and equipment net of proceeds from sales and incentives (Table A.1), the measure that Tables 1 and A.4 use. *Alphabet:* 10-K for 2025 (filed February 5, 2026), risk factors; press release of June 1, 2026 (Form 8-K of June 4, 2026, Exhibit 99.1); prospectus supplement (Form 424B5) dated June 2, 2026 and filed June 4, 2026, Use of Proceeds; 10-Q for the quarter ended June 30, 2026 (filed July 23, 2026). The 2025 range is from the earnings release of October 29, 2025. *Microsoft:* 10-K for fiscal 2026 (filed July 29, 2026), risk factors, liquidity section, and leases note; presentation furnished with Form 8-K of September 2, 2026 (Exhibit 99.1); earnings call of October 29, 2025, for the July–September 2025 figure, which, like the July–September 2026 outlook, includes finance leases. *Meta:* earnings releases of April 29 and July 29, 2026 (Form 8-K, Exhibit 99.1); 10-Q for the quarter ended June 30, 2026 (filed July 30, 2026), risk factors. The 2025 range is from the earnings release of October 29, 2025. *Oracle:* 10-K for fiscal 2026 (filed June 22, 2026), liquidity section and risk factors; earnings release of September 10, 2026 (Form 8-K, Exhibit 99.1); 10-Q for the quarter ended August 31, 2026 (filed September 11, 2026), liquidity section. The year-earlier remaining performance obligations are from the earnings release of September 9, 2025. *CoreWeave:* earnings release of May 7, 2026 (Form 8-K, Exhibit 99.1); 10-Q for the quarter ended March 31, 2026 (filed May 8, 2026), risk factors. *IREN:* business update and results release of May 7, 2026 (Form 8-K, Exhibit 99.1); 10-Q for the quarter ended March 31, 2026 (filed May 8, 2026), management's discussion and analysis. *SpaceX:* Form S-1 filed May 20, 2026, management's discussion and analysis and risk factors. AI-segment capital expenditures are from the segment note of the 10-Q for the quarter ended June 30, 2026.

Table A.5: Assumed and projected 2027 capital expenditure ($ billions)

	Assumed: same net addition as 2026 (2022 dollars)				Projected (current dollars)			
	2026 capex	Net addition in 2026	Depreciation in 2027	2027 used	Latest quarter	Quarter	Growth (%)	2027
Amazon	176	103	99	202	53.1	April–June	69	388
Alphabet	163	116	76	192	44.9	April–June	100	400
Microsoft	154	97	82	179	41.0	April–June	69	300
Meta	122	81	61	142	31.1	April–June	83	248
Oracle	81	67	31	98	28.5	June–August	235	348
xAI	49	44	15	60	—			—
CoreWeave	24	20	9	29	—			—
IREN	6	6	2	8	—			—
Applied Digital	5	5	2	7	—			—
Hut 8	2	2	0	2	—			—
Total	782	540	377	917				

Notes: Each firm's capital stock is built from its whole-billion entries in Table 1 for 2022–2026, converted to 2022 dollars with the consumer price index, as in Appendix A.4, starting from $K_{i0} = I_{i,2022}/\delta_a$ with $\delta_a = 0.25$. Net addition in 2026 is 2026 capital expenditure less depreciation on the stock entering 2026. The 2027 value used is the net addition in 2026 plus depreciation on the stock entering 2027, which equals 2026 capital expenditure plus depreciation on the net addition in 2026. The assumed values are in 2022 dollars. Table 1 shows the 2027 value used in 2026 dollars: the value here before rounding, multiplied by $331.837/292.626 = 1.134$, the ratio of the 2026 to the 2022 consumer price index, and rounded to the nearest billion. Latest quarter is capital expenditure in the most recent quarter reported, measured as in Table 1 for 2022–2025, and growth is relative to the same quarter a year earlier, from the firms' earnings releases and, for Microsoft, its earnings calls. The projected value, in current dollars, grows spending from the latest quarter at the quarterly rate $(1+g)^{1/4} - 1$, where g is that growth, and sums the four quarters that make up 2027 in Table 1. It is not computed for the five other firms. Entries are rounded separately, so rows and columns may not add exactly.

Table A.7: Leases not yet commenced ($ billions, current dollars)

	As of	Amount	Earlier (as of)	Commencement; lease terms
Amazon	Jun 2026	137.2	96.4 (Dec 2025)	Not stated; 11.7 due in 2027
Alphabet	Jun 2026	85.2	58.5 (Dec 2025)	2026–2031; 1–26 years
Microsoft	Jun 2026	329.1	92.7 (Jun 2025)	FY2027–FY2033; 1–20 years
Meta	Jun 2026	279.0	103.8 (Dec 2025) 34.1 (Dec 2024)	2026–2036; over 1 to 30 years
Oracle	Aug 2026	288	260 (May 2026) 43.4 (May 2025)	Q2 FY2027–FY2029; 15–19 years
CoreWeave	Mar 2026	40.7	38.5 (Dec 2025)	2026–2029; 5–16 years
SpaceX	Dec 2025	1.6	n.a.	Not stated; average 7.2 years

Notes: Amounts are the future payments that the firm discloses under leases signed but not yet commenced at the date shown. Dates are month ends. FY is the firm's fiscal year (Appendix A.1). Amounts are rounded to one decimal; those without a decimal are filed in whole billions. *Amazon:* Form 10-Q for the quarter ended June 30, 2026 (filed July 31, 2026), Note 4, line "Leases not yet commenced" of the commitments table, and Form 10-K for 2025 (filed February 6, 2026), Note 7. Amazon does not state commencement dates or lease terms. *Alphabet:* Form 10-Q for the quarter ended June 30, 2026 (filed July 23, 2026), Note 4 (leases primarily related to data centers), and Form 10-K for 2025 (filed February 5, 2026). The 10-K reports short-term and long-term future lease payments of $5.8 and $52.7 billion; the table gives their sum. Note 4 of the 10-Q also reports a short-term lease signed in June 2026, with a non-cancelable commitment of about $5.8 billion, to commence in the third quarter of 2026. *Microsoft:* Forms 10-K for fiscal 2026 (filed July 29, 2026) and fiscal 2025 (filed July 30, 2025), Note 13, leases primarily for datacenters. The fiscal 2026 filing states that some of these leases are subject to contractual conditions being met. The Form 10-K for fiscal 2024 reports $8.6 billion of operating leases and $108.4 billion of finance leases not yet commenced at June 30, 2024. *Meta:* Form 10-Q for the quarter ended June 30, 2026 (filed July 30, 2026), Note 9 (leases for data centers, colocations, and certain network infrastructure); Forms 10-K for 2025 (filed January 29, 2026) and 2024 (filed January 30, 2025). The 10-Q also reports data center leases of about $68 billion signed in July 2026, expected to commence in 2027 and 2028, with terms of 18 to 20 years. *Oracle:* Form 10-Q for the quarter ended August 31, 2026 (filed September 11, 2026); Forms 10-K for fiscal 2026 (filed June 22, 2026) and fiscal 2025 (filed June 18, 2025). These leases are substantially all related to data center arrangements. *CoreWeave:* Form 10-Q for the quarter ended March 31, 2026 (filed May 8, 2026) and Form 10-K for 2025 (filed March 2, 2026), leases primarily for data centers and office buildings. The amounts exclude leases whose payments depend on the lessor's construction costs. As of March 31, 2026, one of these gives access to 525 MW of power, with total rent of $18.7 to $19.6 billion over sixteen years. *SpaceX:* Form S-1 (filed May 20, 2026), operating leases signed but not commenced at December 31, 2025. The amount is for SpaceX as a whole, including its AI segment (xAI). The Form 10-Q for the quarter ended June 30, 2026 reports no material changes in the lease portfolio since December 31, 2025. IREN reports no amount for leases not yet commenced.

Table A.8: Multi-year commitments related to future capacity ($ billions, current dollars)

Firm	Item	As of	Amount	Year earlier	Detail
Amazon	Unconditional purchase obligations	Jun. 30, 2026	130.1	n.a.	33.0 due in 2027. Mar. 31, 2026: 103.8, with 15.8 due in 2027. Dec. 31, 2025: 84.8, with 8.9 due in 2027.
	Customer commitments not yet recognized, primarily AWS	Jun. 30, 2026	496	n.a.	Mar. 31, 2026: 364; Dec. 31, 2025: 244. Weighted-average remaining life 6.4 years (4.1 years at Dec. 31, 2025).
	OpenAI commitment, expansion	Q1 2026	100.0	n.a.	Over 8 years, adding to an existing 38.0. About 2 GW of Trainium capacity, ramping from 2027.
	Anthropic commitment, expansion	Q2 2026	>100.0	n.a.	Over 10 years. Up to 5 GW of Trainium chips.
Alphabet	Purchase commitments and other contractual obligations	Jun. 30, 2026	811.0	n.a.	200.7 short-term. Mar. 31, 2026: 332.4; Dec. 31, 2025: 149.1. Note 10: 707.0 fixed or guaranteed under contracts with more than one year to run, mostly long-term supply agreements to secure future production capacity, generally fulfilled through 2030 (Mar. 31, 2026: 232.7).
	Revenue backlog	Jun. 30, 2026	519.5	n.a.	513.9 for Google Cloud; just over half within 24 months. Mar. 31, 2026: 467.6; Dec. 31, 2025: 242.8. Includes TPU supply agreements ("multiple gigawatts" at Mar. 31, 2026), with the significant majority of revenue in 2027.
Microsoft	Purchase commitments	Jun. 30, 2026	194.1	110.0	169.0 due in FY2027; a year earlier, 103.9 due in FY2026. Primarily datacenters.
	Construction commitments	Jun. 30, 2026	34.6	32.1	29.8 due in FY2027; a year earlier, 26.9 due in FY2026. Primarily datacenters.

Table A.8, continued

Firm	Item	As of	Amount	Year earlier	Detail
	Commercial remaining performance obligation	Jun. 30, 2026	678	368	About 30% within 12 months; a year earlier, about 40% of total remaining performance obligation. Weighted-average duration about 2.3 years.
Meta	Non-cancelable contractual commitments	Jun. 30, 2026	349.3	n.a.	81.7 due in 2027. Mar. 31, 2026: 237.7, with 47.7 due in 2027. Dec. 31, 2025: 131.0, with 22.2 due in 2027. In addition, contingent obligations to buy up to 14.7 of cloud capacity over five years.
Oracle	Remaining performance obligations	Aug. 31, 2026	664	455	About 13% within 12 months and 37% in months 13–36. May 31, 2026: 638.
	Unconditional purchase and other obligations	Aug. 31, 2026	34.2	n.a.	5.4 due by May 2027 and 6.2 in FY2028. May 31, 2026: 13.3, plus 19 entered after that date for cloud infrastructure assets over five years. May 31, 2025: 2.2.
CoreWeave	Remaining performance obligations	Mar. 31, 2026	98.8	n.a.	36% within 24 months. Dec. 31, 2025: 60.7; Dec. 31, 2024: 15.1.
	Contracted power	Mar. 31, 2026	>3.5 GW	n.a.	Up more than 400 MW in the quarter; about 3.1 GW at Dec. 31, 2025. Active power above 1 GW.
IREN	Commitments	Mar. 31, 2026	11.9	n.a.	Nearly all due within 12 months, that is, by March 2027. Jun. 30, 2025: 0.4.
	GPU-services contract with Microsoft	Nov. 2, 2025	9.7	n.a.	About 5 years; 20% of contract value paid before each tranche.
	GPU-services contract with NVIDIA	May 7, 2026	3.4	n.a.	Five-year term, through 2032; three tranches targeted for deployment during 2027.
SpaceX	Non-cancelable contractual commitments	Jun. 30, 2026	28.0	n.a.	22.2 due in 2027. Dec. 31, 2025: 25.5, with 21.5 due in 2027. Includes the EchoStar spectrum purchase (about 19.6), expected to close about Nov. 30, 2027.

Table A.8, continued

Firm	Item	As of	Amount	Year earlier	Detail
	Cloud services agreements (AI segment)	Q2 2026	14.1	n.a.	Contracted sales for access to compute capacity. Company-wide backlog: 47.5 at Jun. 30, 2026; 27.6 at Mar. 31, 2026.

Notes: Amounts are rounded to one decimal; those without a decimal are filed in whole billions. "Year earlier" is the same item one year before the as-of date; n.a. means that no filing reports it, and the Detail column then gives any earlier values. For individual contracts, "As of" gives the quarter or date of the announcement or signing. "Due in 2027" refers to calendar 2027, and FY to the firm's fiscal year (Appendix A.1). Remaining performance obligations, revenue backlog, and contracted sales are revenue that the firm has contracted with customers but not yet recognized. From the first quarter of 2026, Alphabet's revenue backlog also includes contracts with an original term of one year or less (about 7.3 at March 31, 2026), so the December 31, 2025 value is on a narrower basis. Amazon's OpenAI and Anthropic amounts include obligations related to the performance of AWS chips. IREN counts its Microsoft contract in remaining performance obligations only as each tranche is delivered and accepted. Apart from the EchoStar purchase, SpaceX's commitments are primarily for AI infrastructure and third-party cloud capacity. About 11.1 of the EchoStar purchase price is payable in stock at the closing. The filing does not split the 2027 amount among these uses. Sources: *Amazon:* Form 10-Q for the quarter ended June 30, 2026 (filed July 31, 2026), Notes 1 and 4; Form 10-Q for the quarter ended March 31, 2026 (filed April 30, 2026); Form 10-K for 2025 (filed February 6, 2026); earnings release of April 29, 2026 (Form 8-K, Exhibit 99.1), for the gigawatt figures. *Alphabet:* Form 10-Q for the quarter ended June 30, 2026 (filed July 23, 2026), Notes 2 and 10 and management's discussion and analysis; Form 10-Q for the quarter ended March 31, 2026 (filed April 30, 2026); Form 10-K for 2025 (filed February 5, 2026). *Microsoft:* Form 10-K for fiscal 2026 (filed July 29, 2026), management's discussion and analysis (contractual obligations) and Notes 6 and 12; Form 10-K for fiscal 2025 (filed July 30, 2025). *Meta:* Form 10-Q for the quarter ended June 30, 2026 (filed July 30, 2026), Note 9; Form 10-Q for the quarter ended March 31, 2026 (filed April 30, 2026); Form 10-K for 2025 (filed January 29, 2026). *Oracle:* Form 10-Q for the quarter ended August 31, 2026 (filed September 11, 2026), Notes 1 and 6 and management's discussion and analysis; Form 10-K for fiscal 2026 (filed June 22, 2026); Form 10-K for fiscal 2025 (filed June 18, 2025). *CoreWeave:* Form 10-Q for the quarter ended March 31, 2026 (filed May 8, 2026), Note 2; Form 8-K of May 7, 2026, Exhibit 99.1; Form 10-K for 2025 (filed March 2, 2026). *IREN:* Form 10-Q for the quarter ended March 31, 2026 (filed May 8, 2026), Notes 5, 21, and 22, and Part II. *SpaceX:* Form 10-Q for the quarter ended June 30, 2026 (signed August 4, 2026), Notes 2, 6, and 16; Form S-1 (filed May 20, 2026), Notes 6 and 17; second-quarter 2026 earnings release (Form 8-K, Exhibit 99.1).

A.3 Capital expenditure of the other five firms

No double counting. CoreWeave's 10-K for 2025 (Note 1) states that its customer arrangements do not meet the definition of a lease and are accounted for as service contracts.

A hyperscaler that buys computing capacity from CoreWeave therefore records the payments as an operating expense, and the equipment stays on CoreWeave's balance sheet. Adding CoreWeave's capital expenditure to that of the hyperscalers does not count the same equipment twice. IREN's contracts with Microsoft and NVIDIA are likewise contracts to provide dedicated GPU services.

Applied Digital's Polaris Forge 1 campus in North Dakota is leased to CoreWeave, and Hut 8's River Bend campus to Fluidstack, with a backstop from Google for the rent; the two firms' other campuses are leased to technology companies, including hyperscalers, that the filings describe as investment grade without naming them. Oracle delivers data-center capacity to Amazon, Alphabet, and Microsoft under its MultiCloud partnership. The five largest firms do not count assets acquired under leases, except that Meta adds principal payments on finance leases and Microsoft includes assets acquired under finance leases (Appendix A.1). CoreWeave leases all of its data centers and counts the property and equipment it pays for, and xAI's AI-segment line is cash purchases of property and equipment. The leased campuses and the MultiCloud capacity therefore enter Table 1 through their owners. The only possible overlaps are through Meta's principal payments, about 3.5% of its capital expenditure in 2025 (Table A.1), and through Microsoft's finance leases.

Capital expenditure by firm. Table A.9 reports the capital expenditure of xAI, CoreWeave, IREN, Applied Digital, and Hut 8. SpaceX acquired xAI in February 2026 and reports it as its AI segment (Space Exploration Technologies Corp., 2026). The entries before 2024 are small or absent because IREN, Applied Digital, and Hut 8 then served bitcoin mining, xAI began operations only in March 2023, and CoreWeave was small.

Sources by firm. For CoreWeave, 2023–2025 capital expenditure is from the cash-flow statement in the 10-K for 2025 (filed March 2, 2026, accession 0001769628-26-000104), and

Table A.9: Capital expenditure of the other firms in Table 1 ($ billions, current dollars)

	2022	2023	2024	2025	2026 (est.)	2027 (used)	Source
xAI[§]	—	0.5	5.6	12.7	55.2	68	SpaceX S-1 and 10-Q (AI segment)
CoreWeave[†]	—	2.9	8.7	10.3	27.0	33	10-K and 10-Q (cash)
IREN[‡]	—	—	—	1.8	7.3	9	10-K, 10-Q, and 6-K
Applied Digital[¶]	—	0.1	0.5	1.5	6.5	7	10-K and 10-Q (AI segments)
Hut 8[‖]	—	—	—	—	1.8	3	10-Q (2026 construction)
Subtotal	**0**	**3**	**16**	**26**	**97**	**120**	

Notes: "—" indicates no or negligible AI-related capital expenditure. The 2027 values are those of Table 1, in 2026 dollars; Table A.5 reports them in 2022 dollars. The subtotal row sums the whole-billion entries of Table 1, so it can differ from the sum of the one-decimal figures above. [†]CoreWeave's figures are its cash purchases of property and equipment, the measure its 10-K reports for its capital investments. Its statement of cash flows also reports, as a non-cash item, property and equipment added on credit ($11.2B in 2025), which it pays for later, through the cash purchase line or, for equipment financed by its suppliers, through repayments of debt; adding the two lines would count purchases on credit twice. The 2026 figure is cash purchases in the first half, $14.1B, plus twice the second quarter, $6.4B. CoreWeave's 2026 guidance of $31–35B, given on its May 7, 2026 earnings call, is reported in the press, and its basis is not stated. Counting equipment when received, whether paid for in cash, on credit, or financed by suppliers (3.4, 9.1, and 20.6 for 2023–2025, and the guidance midpoint of 33 for 2026), the boom size would be 3.31 rather than 3.28 and p_* 0.034 rather than 0.035. [‡]IREN's fiscal year ends June 30. IREN pays for hardware in advance and reports these payments on a separate line, "Payments for computer hardware"; the figures add this line to "Payments for property, plant and equipment, net of computer hardware." Calendar years are built from its annual, half-year, and quarterly cash-flow statements. IREN does not separate mining from AI hardware. Its hardware payments in fiscal 2024 were for mining hardware, so its 2024 spending, $1.0B, is excluded, and its 2025 entry includes some mining hardware. The 2026 figure is the first half, $3.3B, plus twice the second quarter, $2.0B. [§]The figures are the capital expenditure of SpaceX's AI segment, which equals the segment's cash purchases of property and equipment. The 2026 figure is capital expenditure in the first half, $23.6B, plus twice the second quarter, $15.8B. [¶]Applied Digital's figures are the capital expenditure that its segment note reports for its HPC Hosting and Cloud Services segments, which from fiscal 2025 on account for nearly all of it. Its fiscal year ends May 31, and its figures are built in the same way as Oracle's: the figure for year T covers December of year $T-1$ through November of year T, with each fiscal year's segment capital expenditure divided between calendar years in proportion to its cash purchases of property and equipment. September–November 2026 is set equal to June–August 2026, $2.1B of cash purchases, scaled by fiscal 2026's ratio of segment capital expenditure to cash purchases, 1.04. [‖]Hut 8's capital expenditure before 2026 was mainly for bitcoin mining, which we exclude. In 2026 it is the construction of two AI data centers, River Bend and Beacon Point. The 2026 figure is capital expenditure in the first half, $0.62B, plus twice the second quarter, $0.58B.

2026 from the Forms 10-Q for the quarters ended March 31 and June 30, 2026 (the latter filed August 12, 2026, accession 0001769628-26-000366); the additions on credit are the line "Liabilities related to property and equipment additions, including OEM financed additions" in the same statements. For IREN, the cash-flow statements for fiscal 2024 and 2025 are from the 10-K for fiscal 2025 (filed August 28, 2025), those for the half-years ended December 31, 2023 and 2024 from its Form 6-K of February 12, 2025 (accession 0001878848-25-000009; prepared under IFRS), that for the nine months ended March 31, 2026 from the 10-Q filed May 8, 2026, and that for fiscal 2026 from the 10-K filed August 27, 2026 (accession 0001878848-26-000052); the January–March 2026 figures are from its results release of May 7, 2026. For xAI, the AI segment's capital expenditure for 2023–2025 is from the segment note to the financial statements in SpaceX's S-1, filed May 20, 2026, accession 0001628280-26-036936, and for 2026 from SpaceX's Form 10-Q for the quarter ended June 30, 2026.[10]

For Applied Digital, fiscal-year cash purchases of property and equipment and segment capital expenditure are from the Forms 10-K for fiscal 2022–2026; the half-years ending in November and the quarter ended August 31, 2026 are from the Forms 10-Q. For Hut 8, the 2026 figures are from the Forms 10-Q for the quarters ended March 31 and June 30, 2026.

Private neoclouds. Table 1 does not include the private neocloud providers Crusoe and Lambda, because neither publishes financial statements.

[10]The AI segment is SpaceX's AI compute, Grok, and X, the former Twitter, which xAI acquired in March 2025. SpaceX accounts for the mergers as reorganizations of entities under common control, so the segment's figures include X in every year. Twitter's capital expenditure was $1.0 billion in 2021 (Form 10-K for 2021), against $55 billion for the AI segment in 2026. SpaceX's other two segments are Space and Connectivity (Starlink).

A.4 Capital stock entering 2028

Investment in year t is the calibration total in Table 1 in 2022 dollars, and it is in place at the start of year $t+1$, as in the law of motion (2). Because 2022 investment replaces depreciation, the initial capital stock at the start of 2023 is $K_{a0} = 156/0.25 = 624$ (billions of 2022 dollars). With $\delta_a = 0.25$ and the totals rounded to one decimal, the capital stock at the start of each following year is

$$\begin{aligned}
&\text{start of 2024:} && 0.75 \times 624 + 152.7 = 620.7,\\
&\text{start of 2025:} && 0.75 \times 620.7 + 254.7 = 720.2,\\
&\text{start of 2026:} && 0.75 \times 720.2 + 427.2 = 967.3,\\
&\text{start of 2027:} && 0.75 \times 967.3 + 782.2 = 1{,}507.7,\\
&\text{start of 2028:} && 0.75 \times 1{,}507.7 + 917.3 \approx 2{,}048.0.
\end{aligned}$$

For the five largest firms alone, whose totals in 2022 dollars for 2023 through 2027 are 149.8, 239.7, 403.5, 696.7, and 812.6, the same steps give 617.8, 703.1, 930.9, 1,394.8, and 1,858.7 at the start of 2024 through 2028.

Time to build. The calculation treats each year's capital expenditure as capital in service at the start of the next year. Part of the spending is on data centers that are still under construction. Meta reports construction in progress, mostly data centers, network infrastructure and servers, of \$26.8B at the end of 2024, \$50.5B at the end of 2025, and \$61.0B at the end of March 2026 (10-K for 2025 and 10-Q for the first quarter of 2026). Spending that is not yet in service by 2028 overstates the productive capital stock entering 2028. If the firms are building toward capacity that comes into service later, however, capital reaches its target after the end of 2027, which raises the boom size implied by the same spending. We do not model time to build, and the net effect on the boom size is ambiguous.

B Labor

The production function $Y_j = z_j^{1-\alpha} K_j^{\alpha}$ can be interpreted as Cobb-Douglas in capital and labor, $Y_j = A_j K_j^{\alpha} L_j^{1-\alpha}$, where A_j is total factor productivity, the sector's labor supply is fixed at $\bar{L}_j$, and $z_j = A_j^{1/(1-\alpha)} \bar{L}_j$. The firm's problem (4) can then be viewed in two ways. The first, which we adopt in the text, is an optimization problem with no labor: z_{at}, and with it $\bar{L}_a$, is fixed. The second is an optimization problem at a fixed wage. A firm with capital K takes the wage w as given and decides it can hire as much labor as it wants at that wage; it does not internalize the externality it imposes on the sector through the wage. Writing $F(K, L) = A_a K^{\alpha} L^{1-\alpha}$, its profit is

$$\pi(K; w) = \max_L \left[F(K, L) - wL \right] = \phi(w)\, K, \qquad \phi(w) = \alpha A_a^{1/\alpha} \left(\frac{1-\alpha}{w} \right)^{(1-\alpha)/\alpha},$$

which is linear in the firm's capital, with optimal labor $L^*(K, w) = K \left((1-\alpha) A_a / w \right)^{1/\alpha}$. Differentiating with respect to K at the fixed wage, with L^* responding to K,

$$\frac{\partial \pi}{\partial K} = F_K + (F_L - w) \frac{\partial L^*}{\partial K} = F_K,$$

where the second term is zero by the first-order condition for labor (the envelope theorem), and there is no term in $\partial w / \partial K$ because the firm takes the wage as given. Freely adjusting labor thus leads to a first-order condition for investment, $\mathbb{E}_t[\phi(w_{a,t+1})] = r_a + \delta_a$, in which the firm's capital does not appear. Capital comes back in through the sector-wide wage. Summing labor demand over firms, with sector capital K_a and labor supply $\bar{L}_a$, the labor market clears at $w_a = (1-\alpha) A_a (K_a / \bar{L}_a)^{\alpha}$, so that $\phi(w_a) = \alpha (z_a / K_a)^{1-\alpha}$, the marginal product of sector capital, giving the same first-order condition as the original problem, (5). At this wage the wage bill $w_a \bar{L}_a$ equals $(1-\alpha) Y_a$ and earnings net of wages equal αY_a. Investment is therefore the same whether earnings are measured gross or net of wages, while

the level of V_{at} is lower by the present value of the wage bill.

C AI-capital depreciation and the required return

The baseline rate $\delta_a = 0.25$ captures the relatively rapid depreciation of AI investment, much of which is computing equipment. Under the geometric law of motion (2), 32% of a cohort remains after four years and 18% after six. The non-AI depreciation rate, by contrast, is $\delta_n = 0.07$.

The largest firms' accounting policies use straight-line depreciation. Between 2022 and 2025, four of them—Microsoft, Alphabet, Meta, and Amazon—extended their estimated useful lives for server and networking equipment: Microsoft and Alphabet to six years (Alphabet Inc., 2023, Microsoft Corporation, 2023), Meta to 5.5 years for certain server and network assets (Meta Platforms, Inc., 2025), and Amazon in stages to six years (Amazon.com, Inc., 2023, 2025). Amazon partially reversed the extension in January 2025, shortening the lives of a subset of its servers and networking equipment from six to five years, citing "an increased pace of technology development, particularly in the area of artificial intelligence and machine learning" (Amazon.com, Inc., 2025). These useful lives do not translate directly into geometric depreciation rates. Straight-line depreciation writes an asset off completely over a fixed horizon, whereas geometric depreciation leaves a fraction of each cohort at every horizon. We consider $\delta_a = 0.15$ as an alternative with more durable capital: 38% of a cohort remains after six years at this rate.

Depreciation affects the inferred boom size through two opposing channels. Recall that (12) identifies the boom size from $K_{a,2028}/K_{a0}$, where the initial stock entering 2023 is $K_{a0} = I_{a0}/\delta_a$. To see the first channel, temporarily set aside depreciation during the investment period, while retaining this initial stock. All investment then adds to the stock, without replacing any capital lost during the buildout. Denoting the resulting hypothetical stock by

$\widetilde{K}_{a,2028}$, equation (11) simplifies to

$$\frac{\widetilde{K}_{a,2028}}{K_{a0}} = 1 + \frac{\sum_{t=2023}^{2027} I_{at}}{K_{a0}}.$$

Investment adds dollar amounts to the capital stock. A lower depreciation rate implies a larger initial stock, so the same investment represents a smaller proportional expansion. This channel therefore implies a smaller productivity boom at a lower depreciation rate.

A second effect pushes in the opposite direction. In the actual capital accumulation equation, part of investment replaces depreciation of the initial stock and of capital installed during the buildout. The more investment goes toward replacement, the less remains to expand the stock, and the smaller the boom needed to justify the observed expenditure. Conversely, at a given initial stock, lower depreciation reduces this replacement requirement and leaves more capital in place in 2028, implying a larger boom.

Restoring depreciation during the buildout and substituting $K_{a0} = I_{a0}/\delta_a$ into (11) gives the exact expression

$$e^{\hat{\xi}} = \frac{K_{a,2028}}{K_{a0}} = (1-\delta_a)^5 + \delta_a \sum_{t=2023}^{2027} (1-\delta_a)^{2027-t} \frac{I_{at}}{I_{a0}}. \tag{14}$$

With the calibration totals of Table 1 in 2022 dollars through 2026, and with the 2027 value recomputed at each rate to keep the net addition to the stock at its 2026 level ($868B at $\delta_a = 0.15$ against $917B at baseline, in 2022 dollars), the first effect dominates. Lowering δ_a from 0.25 to 0.15 raises the initial stock from $624B to $1,040B, an increase of 67%. The capital stock entering 2028 also rises, from $2,048B to $2,539B, but by only 24%. Its ratio to the initial stock therefore falls: $e^{\hat{\xi}}$ is 2.44 rather than 3.28.

Table C.1 compares the baseline (Panel A) with a lower depreciation rate (Panel B) and the industry hurdle rate (Panel C), including the AI capital stock alongside the quantities

in Table 2. In Panels A and B, the boom size is inferred from the same investment data through 2026 using that panel's depreciation rate. In each panel, p_* is at its baseline value of 0.035, which is below the bound $r_a/(e^{\hat{\xi}}-1)$ given by that panel's boom size and required return.

Table C.1: Post-window output shares and GDP growth under alternative depreciation and required-return calibrations

	AI share (%)	Cum. GDP growth (%)	AI capital ($B)	Post-window I_a ($B/yr)
Panel A. Baseline: $\delta_a = 0.25$, $e^{\hat{\xi}} = 3.28$, $r_a = 0.08$, $p_* = 0.035$				
Initial	2.4	—	624	156
Moderate	7.5	5.6	2,179	545
Transformative	21.1	23.7	7,153	1,788
Singularity	46.7	83.2	23,476	5,869
Panel B. $\delta_a = 0.15$, *re-estimated* $e^{\hat{\xi}} = 2.44$, $p_* = 0.035$				
Initial	2.8	—	1,040	156
Moderate	6.6	4.1	2,648	397
Transformative	14.6	13.9	6,466	970
Singularity	29.5	37.9	15,788	2,368
Panel C. $r_a = 0.15$, $\delta_a = 0.25$, $e^{\hat{\xi}} = 3.28$, $p_* = 0.035$				
Initial	2.9	—	624	156
Moderate	9.0	6.8	2,179	545
Transformative	24.5	28.7	7,153	1,788
Singularity	51.6	100.8	23,476	5,869

Notes: Construction parallels Table 2; dollar amounts are in 2022 dollars. "AI capital" is the initial stock $K_{a0} = I_{a0}/\delta_a$ in the initial row and the post-window steady-state stock K_a^* otherwise; post-window $I_a = \delta_a K_a^*$. Cumulative GDP growth is the one-time increase in the level of GDP from the initial economy to the post-window steady state. Initial investment is the 2022 capital expenditure of $156B; δ_n is at its baseline value in all panels, r_n is set in each panel so that initial GDP equals measured 2022 GDP, and r_a is at its baseline value of 0.08 in Panels A and B. In every panel $p_* = 0.035$, below that panel's bound $r_a/(e^{\hat{\xi}}-1)$: 0.0351, 0.0555, and 0.0657 in Panels A, B, and C. Calculations use the unrounded value of $\hat{\xi}$.

Panel B uses $\delta_a = 0.15$ throughout, with the boom size implied by (14): $e^{\hat{\xi}} = 2.44$, or $\hat{\xi} = 0.89$. With p_* at its baseline value, the smaller boom lowers expected productivity growth to 5.0% per year, compared with 8.0% at baseline. Output shares and GDP gains are lower

than at baseline in all three scenarios, although in the moderate scenario capital is higher: the post-window stock is \$2,648B against \$2,179B, while output is \$1,777B against \$2,070B. Output per unit of capital in the post-window steady state, $Y_a/K_a^* = (r_a + \delta_a)/(\alpha\,\Omega(p_*))$, is 0.67 in Panel B against 0.95 in Panel A, and the gain in output is measured from an initial level of \$718B rather than \$618B. In the transformative and singularity scenarios, with two and three booms in all, the smaller boom compounds ($e^{2\hat{\xi}}$ is 6.0 rather than 10.8 and $e^{3\hat{\xi}}$ is 14.6 rather than 35.4), and capital is lower than at baseline as well. With booms arriving from 2031 on at the rate p_*, as in Figure 2, expected cumulative GDP growth by 2050 is 15%, 39%, and 100% in the three scenarios, against 32%, 109%, and 363% at baseline. The lower GDP gains reflect the smaller boom: smaller productivity gains during the window and, with p_* unchanged, slower growth after it.

The required return. The required return r_a does not affect the boom size, because the ratio in Proposition 3.1 cancels it. It enters the growth calculations through the initial productivity level z_{a0}, which is set so that K_{a0} satisfies the optimality condition (8). At baseline it also sets p_*, through the bound in (15) below. At a given p_*, the post-window capital stock after n booms, $K_{a0}e^{n\xi}\,\Omega(p_*)^{1/(1-\alpha)}$, does not depend on r_a, and output per unit of capital, $(r_a + \delta_a)/(\alpha\,\Omega(p_*))$, rises with it.

For the firm's problem to have a finite solution, the expected gross growth factor of productivity must be below the gross discount rate:

$$1 + r_a > (1 - p) + p\,e^{\xi}. \tag{15}$$

Given that p is subject to unanticipated structural breaks, (15) may be violated for a value of p that obtains for a finite number of periods. We set p_* just below the value at which (15) holds with equality, $r_a/(e^{\hat{\xi}} - 1) = 0.0351$ at baseline, so that $p_* = 0.035$, and hold it at this

value in Panels B and C.

Panel C of Table C.1 sets r_a to the 15% hurdle rate of Vanguard Group (2025), holding δ_a, the boom size, and p_* at their baseline values. Output per unit of capital rises by 21%, before and after the window: initial AI output is \$749B against \$618B at baseline, and the initial AI share is 2.9% against 2.4%. The post-window capital stocks are as at baseline. In the moderate scenario, GDP in the post-window steady state is 6.8% above its initial value, against 5.6% at baseline, and with booms arriving from 2031 on at the rate p_*, as in Figure 2, expected cumulative GDP growth by 2050 is 38%, 132%, and 440% in the three scenarios, against 32%, 109%, and 363% at baseline. The differences come from the higher output per unit of capital.

D Implications for the equity premium and for interest rates

The model of Sections 3–4 takes the required return r_a as given. This appendix computes the risk-free rate and the equity premium in an endowment economy in which booms of the calibrated size arrive at the rate p_*.

Preferences and consumption. A representative agent has the recursive preferences of Epstein and Zin (1989) with unit elasticity of intertemporal substitution, risk aversion γ, and discount factor β. With unit elasticity the price–consumption ratio is constant, $P_t/C_t = \beta/(1-\beta)$, the return on wealth is $R_{W,t+1} = \beta^{-1} C_{t+1}/C_t$, and the stochastic discount factor is (see Miller et al., 2025)

$$M_{t+1} = \beta \, \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{1-\gamma} \right]^{-1} \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma}. \tag{16}$$

Log consumption growth is

$$\Delta c_{t+1} = \mu + \sigma\,\varepsilon_{t+1} + x_{d,t+1} + x_{b,t+1}, \tag{17}$$

where μ and σ are the mean and volatility of normal-times growth, $\varepsilon_{t+1} \sim N(0,1)$, $x_{d,t+1}$ is a disaster that equals $\xi_d < 0$ with probability p_d and zero otherwise, and $x_{b,t+1}$ is a boom that equals $\xi_b > 0$ with probability p_b and zero otherwise. The three shocks are independent and i.i.d. over time. Following Barro (2006) and Wachter (2013), $p_d = 0.036$ and $\xi_d = -0.15$.

The aggregate boom. A boom multiplies AI-sector productivity by e^{ξ} and AI output by $e^{(1-\alpha)\xi}$ before capital adjusts. We consider the limit in which the AI sector accounts for all output, which obtains after sufficiently many booms, so that a boom raises log aggregate consumption by $\xi_b = (1-\alpha)\xi \approx 0.79$, and set $p_b = p_*$. Using the output response before capital adjusts understates the effect of a boom on consumption in the long run, which is ξ once capital has adjusted.

Asset prices. With

$$\Phi_d(u) = (1-p_d) + p_d\,e^{u\,\xi_d}, \qquad \Phi_b(u) = (1-p_b) + p_b\,e^{u\,\xi_b}, \tag{18}$$

the Euler equation $\mathbb{E}_t[M_{t+1}R_{t+1}] = 1$ gives the log risk-free rate

$$r_f = \rho + \mu + \frac{1-2\gamma}{2}\,\sigma^2 + \log\frac{\Phi_d(1-\gamma)}{\Phi_d(-\gamma)} + \log\frac{\Phi_b(1-\gamma)}{\Phi_b(-\gamma)}, \tag{19}$$

where $\rho = -\log\beta$, and the log equity premium on the consumption claim

$$\log\mathbb{E}[R_W] - r_f = \gamma\,\sigma^2 + \log\frac{\Phi_d(1)\,\Phi_d(-\gamma)}{\Phi_d(1-\gamma)} + \log\frac{\Phi_b(1)\,\Phi_b(-\gamma)}{\Phi_b(1-\gamma)}. \tag{20}$$

For small p_d and p_b the premium is approximately $\gamma\sigma^2 + p_d(e^{\xi_d} - 1)(1 - e^{-\gamma\xi_d}) + p_b(e^{\xi_b} - 1)(1 - e^{-\gamma\xi_b})$: disasters and booms both raise it.

Table D.1 reports r_f and the equity premium without AI booms and with them. At $\gamma = 3$, the equity premium rises from 0.39% to 4.13% and the risk-free rate from 3.63% to 4.03%. The risk-free rate rises because expected consumption growth is higher. The premium rises because the boom is a state in which consumption growth and the return on wealth are both high, which any $\gamma > 0$ penalizes. Higher risk aversion damps the rise in r_f through precautionary saving and amplifies the rise in the premium.

Table D.1: Equilibrium rates of return

	No AI	AI booms
Boom probability p_b	0	0.035
Boom size ξ_b	—	0.792
Panel A: Risk-free rate r_f (%)		
$\gamma = 1$	3.90	5.83
$\gamma = 3$	3.63	4.03
$\gamma = 5$	3.29	3.38
Panel B: Equity risk premium (%)		
$\gamma = 1$	0.12	2.33
$\gamma = 3$	0.39	4.13
$\gamma = 5$	0.72	4.78

Notes: Entries are the equilibrium risk-free rate r_f and the equity risk premium on the claim to aggregate consumption (20), in annual percent, at risk aversion $\gamma \in \{1, 3, 5\}$; the elasticity of intertemporal substitution equals one throughout. The boom size $\xi_b = (1 - \alpha)\xi$ scales the calibrated AI-sector boom $\xi \approx 1.19$ by the output elasticity $1-\alpha = 2/3$; booms arrive with probability $p_b = p_* = 0.035$ per year. Remaining parameters: rate of time preference $\rho = 0.02$; mean consumption growth $\mu = 0.025$ with volatility $\sigma = 0.02$; disasters of size $\xi_d = -0.15$ arrive with probability $p_d = 0.036$.