NBER WORKING PAPER SERIES

CLOUDS OF CREDIT:
MAPPING THE FINANCIAL WEB BEHIND AI COMPUTE

Stijn Van Nieuwerburgh
Zijin Su

Working Paper 36032
http://www.nber.org/papers/w36032

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2026

The authors would like to thank Gaurav Chaudhary and Ziyi Qiu for excellent research assistance. We thank Olivier Darmouni and Arpit Gupta for comments. First draft: November 1, 2026. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 Stijn Van Nieuwerburgh and Zijin Su. 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.

Clouds of Credit: Mapping the Financial Web Behind AI Compute
Stijn Van Nieuwerburgh and Zijin Su
NBER Working Paper No. 36032
November 2026
JEL No. E22, G01, G1, G10, G23, G31, G33, R33

ABSTRACT

We map the financial network supporting the AI buildout. Substantial equity cushions protect creditors against large obligations outstanding in mid-2026, but concentrated exposures and interconnected claims create weak links. Our calibrated model supports a large compute expansion in the coming years. The equilibrium investment exceeds the social optimum because firms do not internalize their effects on competitors' revenues. Debt finances nearly 60% of the expansion, increasing leverage, while non-recourse borrowing leaves project lenders exposed despite equity buffers. In the low-productivity state, losses reach $16.8 trillion across shareholders and creditors, with 40% loss-given-default on debt. Suppressing contractual payment spillovers among the key firms reduces private losses by approximately 30%, revealing substantial network propagation.

Stijn Van Nieuwerburgh
Columbia University
Graduate School of Business
Department of Finance
and NBER
svnieuwe@gsb.columbia.edu

Zijin Su
Columbia University
Graduate School of Business
Finance
ZSu27@gsb.columbia.edu

1 Introduction

The rapid adoption of artificial intelligence across businesses and households rests on a vast physical buildout of chips, data centers, and power infrastructure. Financing this buildout requires trillions of dollars, with projected investment rivaling or exceeding the largest infrastructure booms in U.S. history as a share of GDP (Van Nieuwerburgh, 2026). As the scale of the buildout exceeds what the leading AI and technology companies are able or willing to finance from internal cash flows, they increasingly turn to third-party data-center developers, joint ventures, and external investors. This reliance on outside capital links chipmakers, cloud providers, AI developers, data-center landlords, and lenders through a web of purchase commitments, service contracts, leases, loans, and equity stakes. These arrangements mobilize capital for expansion while creating financial exposures across the firms involved. The revenues supporting one firm's expansion, however, depend on other firms' ability to honor their commitments—and ultimately on the uncertain value that AI creates. Understanding the buildout therefore requires examining the network as a whole. How much compute can this network finance? Will private investment exceed or fall short of the socially optimal scale? And if AI productivity disappoints, how will losses spread through the network, and who will ultimately bear them?

We develop a model in which investment and financing are jointly determined within a network of firms involved in the AI buildout. AI labs and hyperscalers choose how much additional compute to acquire for their own use. We characterize a Nash equilibrium in which each chooses its buildout to maximize the value of its incumbent shareholders' equity, taking the other firms' buildout choices as given and anticipating how its own choice affects contract prices, financing, and repayments. Hyperscalers and neoclouds install and operate the hardware that serves this demand, purchasing chips from a hardware supplier and either owning their facilities or renting them from data-center landlords. Firms finance the buildout with cash, new equity, and loans backed by the cash flows from their investments. Service and rental contracts are priced competitively to compensate suppliers for customer default risk. Once AI productivity is realized, operating income and payments from counterparties determine creditor recoveries and shareholder payoffs throughout the network, subject to contractual payment priorities. The model thus connects the incentives governing investment under uncertainty with the distribution and transmission of losses afterward.

A central contribution of the paper is to construct a detailed, internally consistent account of the financial relationships underlying the AI buildout. These relationships are scattered

across company filings, counterparty disclosures, financing announcements, and private-company documents; no single balance sheet reveals the network. We assemble them into a common ledger at June 30, 2026, identifying who owes whom, the nature and priority of each claim, and the equity stakes that connect firms' shareholder payoffs. We reconcile named bilateral claims with company-level obligations, separating exposures within the network from amounts owed to outside creditors. We include debt, accrued obligations and signed but unfulfilled commitments in the inherited exposures, distinguishing these legacy obligations from new contracts associated with additional investment. Combining this financial ledger with data on owned and rented data-center capacity connects contractual exposures to the physical infrastructure they support. The resulting measurement framework makes the network's structure explicit while distinguishing documented amounts, constructed estimates, and assumptions where disclosures are incomplete.

We use this ledger to calibrate a network comprising NVIDIA, six hyperscalers, two AI labs, two neoclouds, and a data-center landlord sector. The calibration combines measured financial exposures and ownership stakes with assumptions about compute benefits, hardware costs, financing limits, and exposure to AI productivity. A key discipline is that firms' legacy income, new contracts, debt, and equity are valued under the same risk-adjusted pricing system. We jointly calibrate legacy outside-income levels and solve for investment so that the resulting equilibrium matches the reference equity values in our June 30, 2026 ledger, using market capitalizations for listed firms, valuation references for private firms, and an estimated equity value for the landlord sector. This ties the resources supporting the buildout to observed valuations while allowing investment, contract prices, borrowing proceeds, and equity dilution to adjust together.

The baseline equilibrium produces a large buildout: 277.5 GW of additional compute and \$6.30 trillion in capital expenditure (in real, present-value terms). The buildout in the decentralized equilibrium exceeds the unconstrained social optimum of 178.9 GW by 55%. Competitive pricing of service and rental contracts does not eliminate the incentive to overinvest. Each compute user values the revenues generated by its own expansion but does not fully internalize the reduction in other users' revenues as aggregate capacity grows. Financing constraints and network relationships shape where this investment takes place: AI labs generate demand for capacity supplied by hyperscalers and neoclouds, while procurement patterns, ownership stakes, and access to financing influence the allocation across firms. The equilibrium therefore reflects both competition for AI revenues and the network's capacity to finance that competition.

The buildout substantially changes the network's financing structure. The inherited balance sheets are predominantly equity financed, with $1.21 trillion of lender debt against $20.55 trillion of equity. By contrast, the $6.30 trillion expansion draws on $3.73 trillion of new loan proceeds, $2.26 trillion of new equity, and $0.31 trillion of existing cash. Debt therefore finances approximately 59% of new investment, despite substantial equity issuance. Measuring debt as promised repayments discounted at the risk-free rate, the aggregate lender-debt-to-equity ratio rises from 5.9% before the buildout to 26.3% afterward. Service and rental contracts create an additional layer of payment obligations linking firms within the network. Including these contracts and other ordinary obligations, gross contractual obligations rise from 25.4% to 71.2% of aggregate equity value. The expansion thus increases both borrowing from outside lenders and firms' exposure to payments promised by one another.

When uncertainty is resolved, the network determines how the AI productivity shock translates into creditor losses and shareholder payoffs. In the low-productivity state, operating income falls and the AI labs cannot meet their service commitments in full. Their payment shortfalls reduce suppliers' receipts, weakening those firms' ability to pay landlords and lenders and transmitting losses through cross-held equity. Oracle also defaults on its ordinary obligations. New project lenders can suffer losses even when the borrowing firm remains solvent because their recovery is limited to pledged investment cash flows. After consolidating claims within the network, the decline from initial market value borne by outside investors totals $16.87 trillion: approximately $14.64 trillion falls on outside shareholders, $2.07 trillion on outside lenders and $0.16 trillion on other outside creditors.

To isolate the role of contractual payment spillovers within the network of the main AI players, we consider a counterfactual in which the government makes up shortfalls on payments between these firms, holding investment and promised contracts fixed. This backstop requires $5.06 trillion in transfers in the low-AI-productivity state and reduces losses borne by private investors from $16.87 trillion to $11.81 trillion. It prevents missed contractual payments from transmitting losses to counterparties, while leaving equity exposures intact. The exercise shows that contractual payment spillovers materially affect private investors' losses and how those losses are distributed across the network.

2 Related Literature

Our immediate predecessor is Van Nieuwerburgh (2026), who documents the scale of the AI buildout and its transition from internal funding toward external capital and complex financing arrangements. We turn that financial architecture into a measured network in which investment, financing, contract prices, and repayments are jointly determined. Complementary work measures aggregate data-center investment from individual projects (Brandsaas et al., 2025), studies data-center mortgage financing (Gete and Mercader, 2026), and examines project-finance risks (Arun, 2025). Other work studies state investment incentives and local electricity-market effects (Gargano and Giacoletti, 2026; Feher et al., 2026). Our focus is the allocation of capital and risk across the firms financing and using compute. The overlapping investment and commercial relationships documented by Frost et al. (2026) make a consistent account of bilateral exposures particularly important.

The paper connects to research on production networks and financial contagion. Acemoglu et al. (2012) show how input-output relationships transmit firm-level shocks to aggregate activity; Eisenberg and Noe (2001), Elliott et al. (2014), and Acemoglu et al. (2015) characterize payment clearing, cross-ownership, and systemic fragility. Several papers also combine transmission channels. Fischer (2014) prices cross-held equity and liabilities of different seniorities, while Costello (2020) and Altinoglu (2021) connect financing disruptions to supply-chain relationships and trade credit. We bring these insights to a setting in which the same firms are simultaneously customers, suppliers, borrowers, and equity investors. Our contribution combines a detailed ledger of these distinct claims with endogenous buildout, financing constraints, and competitively priced new contracts. This permits us to trace both contractual payment spillovers and the losses ultimately borne by outside investors.

Our investment mechanism relates to corporate finance and the economics of investment booms. Limited liability and inherited debt can separate shareholders' investment incentives from total surplus (Jensen and Meckling, 1976). Our analysis also relates to the financing role of asset ownership: Eisfeldt and Rampini (2009) and Rampini and Viswanathan (2013) show that leasing can relax financing constraints because lessors' stronger repossession rights support greater financing capacity. The priority afforded to our new project lenders on pledged investment cash flows connects to Stulz and Johnson (1985), who show that securing new borrowing on new investment can mitigate debt overhang. Competition for revenues creates a further wedge when firms do not internalize the effect of their expansion on competitors, related to the business-stealing externality in Mankiw and Whinston (1986). Greenwood and

Hanson (2015) connect industry investment waves to expectations that neglect competing capacity. In our Nash equilibrium, firms anticipate competing investment but do not internalize its full revenue consequences. Particularly close is Rungcharoenkitkul (2026), who models an AI investment race with overinvestment, leverage, and contagion through circular equity holdings. We complement that analysis by calibrating a network of service contracts, leases, debt, and equity stakes and jointly determining capacity, financing, and default compensation. The model distinguishes the investment externality from the contractual channels that distribute losses after uncertainty is resolved.

Finally, we connect to asset pricing research on technological change and AI exposure. Pástor and Veronesi (2009) study how uncertainty about technological revolutions affects stock prices. Eisfeldt et al. (2026) measure firms' exposure to generative AI and relate it to equity revaluation; Eisfeldt and Schubert (2025) survey the broader finance literature. We examine how a common AI productivity shock is priced across the claims financing the buildout. Under a common risk-adjusted pricing system, service providers, landlords, and new lenders require compensation for customer default risk, while equity values reflect operating income and counterparty exposures. We take this pricing system as given rather than estimate an AI risk premium. This discipline links the financing available before investment to recoveries and shareholder payoffs afterward, allowing us to assess why fair pricing of new claims need not deliver socially optimal investment.

3 Model

3.1 Participants, Capacity and Timing

Let $\mathcal{I}$ denote the set of firms, modeled explicitly in the network, and use i or j to index its members. Let $\mathcal{A}$ denote the set of AI labs who develop AI models but use other firms' compute capacity, $\mathcal{H}$ the set of hyperscalers, $\mathcal{K}$ the set of neocloud firms, and $\mathcal{B}$ the set of chip producers. Let a, h, k, and b index a lab, hyperscaler, neocloud, and chip builder, respectively. Let L denote the consolidated landlord sector and $\mathcal{L} = \{\mathsf{L}\}$ the set of landlords. Outside shareholders, lenders, and customers complete the economy. Company identities and all numerical assumptions are specified in the calibration Section 4.2.

The economy enters with a given amount of aggregate compute Q^0 and an allocation of that compute across firms q_i^0. Compute is expressed by its power requirements, measured

in gigawatts (GW). Firms also enter with a full set of legacy obligations vis-a-vis other companies inside the network, summarized by the matrix W^0, as well as vis-a-vis outside parties (investors, lenders, other customers and suppliers).

At date 0, firms $i \in \mathcal{A} \cup \mathcal{H}$ choose q_i^1, their additional compute capacity to be delivered in period 1. Let Q^1 denote the total additional capacity, what we will refer to as the size of the AI buildout:

$$Q^1 = \sum_{i \in \mathcal{A} \cup \mathcal{H}} q_i^1, \qquad q_i^1 \geq 0. \tag{1}$$

This buildout is an optimal investment decision made under uncertainty. A state of the world s is realized at time 1 that reveals the productivity of compute capacity, θ_s, capturing its usefulness to society. The (physical) probabilities of the states are denoted by π_s. Let $\bar{\theta} = \sum_s \pi_s \theta_s$ denote the mean physical productivity, which we normalize to $\bar{\theta} = 1$.

Time 0, when legacy obligations are inventoried, is conceived as the end of June 2026. The model condenses the periods after that date into a single future period-1 date.

3.2 Asset Pricing

Let R denote the gross risk-free rate between dates, converting between time-0 and time-1 values. Monetary quantities below are in date-0 present-value units, unless explicitly mentioned. Let $\tilde{\pi}_s$ denote risk-neutral probabilities and M_s the stochastic discount factor (SDF) in state s. Both probability vectors are positive and sum to one. Let X_s be any date-1 payoff and $P(X)$ its date-0 price. Then

$$P(X) = \sum_s \pi_s M_s X_s = \frac{1}{R} \sum_s \tilde{\pi}_s X_s. \tag{2}$$

The relation between risk-neutral and physical probabilities is:

$$\tilde{\pi}_s = R \pi_s M_s, \qquad M_s = \frac{\tilde{\pi}_s}{R \pi_s}. \tag{3}$$

and $\mathbb{E}[M_s] = R^{-1}$. A larger SDF value in a bad aggregate AI state indicates that payments in those states are especially valuable. Investor risk aversion motivates this logic.

3.3 Legacy Assets and Liabilities

The matrix W^0 records legacy liabilities as of time-0, excluding the new obligations taken on as part of the buildout, between the firms in the network $i \in \mathcal{I}$. These obligations take on

three forms that can be ranked in order of priority in bankruptcy: secured (or senior) claims, ordinary (or junior) claims, and equity. Let $W_{ij}^{0,J}$ denote the ordinary-priority amount owed to firm i by issuer j, contracted by the reference date and due at time 1. Rows identify creditors and columns identify obligors. Similarly, $W_{ij}^{0,S}$ are the senior obligations owed to firm i by issuer j. Finally, O_{ij}^0 is the fraction of shares owned by firm i in issuer firm j. Multiplying the ownership fraction O_{ij} by issuer j's book or market equity gives the corresponding dollar value of the holding.

Use a dot to denote summation over the corresponding counterparty index: $W_{i\cdot}^{0,J}$ denotes the ordinary claims to be received by i, an asset of firm i, while $W_{\cdot i}^{0,J}$ denotes the ordinary claims owed by i, a liability of firm i:

$$W_{i\cdot}^{0,J} = \sum_{j \in \mathcal{I}} W_{ij}^{0,J}, \qquad W_{\cdot i}^{0,J} = \sum_{j \in \mathcal{I}} W_{ji}^{0,J}. \tag{4}$$

The same row/column convention applies to senior claims.

Let $D_i^{0,out,S}$ and $D_i^{0,out,J}$ denote issuer i's total inherited senior and junior liabilities to parties outside the modeled network. Let $D_i^{0,S}$ and $D_i^{0,J}$ denote total senior and junior liabilities, summing across amounts owed to firms inside and outside the network.

Let C_i^0 denote inherited cash and H_i^0 other outside financial assets and receivables from firms outside the modeled network $\mathcal{I}$. Let K_i^0 denote the legacy data-center real estate, power infrastructure, and IT hardware owned by firm i. Neoclouds own the IT they operate but lease their facilities, so K_i^0 only includes IT assets for $i \in \mathcal{K}$. Labs $a \in \mathcal{A}$ own neither asset in this model. IT at a leased site belongs to its operator (the tenant), not to the landlord.

The pre-buildout balance sheet of firm i, denoted in book value terms, looks as follows:

$$\begin{aligned} A_i^{0,BV} &= C_i^0 + K_i^0 + H_i^0 + W_{i\cdot}^{0,J} + W_{i\cdot}^{0,S} + \sum_j O_{ij}^0 E_j^{0,BV}, \\ D_i^{0,BV} &= D_i^{0,\text{out},J} + D_i^{0,\text{out},S} + W_{\cdot i}^{0,J} + W_{\cdot i}^{0,S}, \\ E_i^{0,BV} &= A_i^{0,BV} - D_i^{0,BV}. \end{aligned} \tag{5}$$

Legacy obligations include financial debt, accrued payables for IT hardware and compute service contracts, and commenced leases. They also include all leases, purchase, construction and compute commitments signed by the reference date (June 30, 2026), but not yet commenced or fulfilled.[1] These obligations are associated with the legacy business and existing

[1] As we discuss below, long-term lease commitments for data centers are often broken up into a sequence

data center assets, and are distinct from the new contracts generated by the buildout of new data centers.

3.4 The Buildout: Firm Objective and Investment

The AI labs and hyperscalers choose their additional compute to maximize (incumbent) shareholder value. Let $\boldsymbol{q}^1_{-i}$ denote the other firms' additional capacity choices, z a candidate own choice, and $\mathcal{F}$ the set of feasible allocations that admit finite contract prices and satisfy financing, equity-participation and payment-feasibility conditions spelled out below. Let E_i^{MV} denote the market value of the firm's equity after the buildout decision, and e_i the additional equity raised to finance the buildout. Hence, $E_i^{MV} - e_i$ denotes incumbent shareholder value, inclusive of the value creation from the buildout. Each firm $i \in \mathcal{A} \cup \mathcal{H}$ chooses the size of its buildout to maximize shareholder value, holding other firms' choices fixed:

$$q_i^1 = \underset{z \geq 0:\, (z, \boldsymbol{q}^1_{-i}) \in \mathcal{F}}{\arg\max} \{E_i^{MV}(z, \boldsymbol{q}^1_{-i}) - e_i(z, \boldsymbol{q}^1_{-i})\}, \qquad i \in \mathcal{A} \cup \mathcal{H}. \tag{6}$$

subject to the financing and payment-feasibility constraints. As we shall see, the latter depend jointly on all firms' choices.

Firms invest in new compute capacity because it produces future cash flows in period 1. These cash flows are stochastic and depend on the realized productivity of AI in time-1, $\theta(s)$. Denoting the aggregate compute of all other firms except firm i by Q^1_{-i}, we model the benefits of compute in state s as:

$$\mathcal{V}^P(q_i^1; Q^1_{-i}, s) = \theta_s q_i^1 \left(a_V - \frac{b_V}{2}(q_i^1 + Q^1_{-i}) \right). \tag{7}$$

This production function of compute exhibits decreasing returns to scale in the size of the aggregate buildout when $b_V > 0$.

The revenues from new compute are accounted for on the firm's balance sheet as time-1 cash flows from financial assets H, $X_i^{H,1}(s)$:

$$X_i^{H,1}(s) = R\, \mathcal{V}^P(q_i^1; Q^1_{-i}, s) \tag{8}$$

of shorter-term leases. Only the first one in the sequence, the commenced lease, shows up on the balance sheet. The later, uncommenced leases will only show up on the balance sheet when they begin. Since these future leases are associated with existing data centers, we include the uncommenced leases in the date-0 obligations. The same holds for future compute service commitments and future chip purchase agreements, already signed at time zero.

where multiplication by R expresses them in time-1 units.

The model distinguishes between users and owners of compute. Next, we define the investment associated with the buildout, i.e. the capital expenditures undertaken by the owners of compute assets. Since AI Labs $a \in \mathcal{A}$ do not own compute assets, they route their compute requirements through the operators $j \in \mathcal{H} \cup \mathcal{K}$ using compute contracts. Let the procurement shares be denoted by ω_{ja}, with $\sum_j \omega_{ja} = 1$ for each lab a. Hyperscalers $h \in \mathcal{H}$ combine these compute purchases by the labs with their own desired compute expansion q_h^1. They can subcontract a share β_{kh} of their combined compute needs to neoclouds $k \in \mathcal{K}$. Neoclouds differ from hyperscalers in that they do not make a compute investment choice, but rather satisfy others' compute demand. Let d_h denote the total demand handled by h and y_j the new capacity operated by j. Then

$$d_h = \sum_a \omega_{ha} q_a^1 + q_h^1, \qquad y_h = \left(1 - \sum_k \beta_{kh}\right) d_h, \tag{9}$$

$$y_k = \sum_a \omega_{ka} q_a^1 + \sum_h \beta_{kh} d_h, \qquad \sum_{j \in \mathcal{H} \cup \mathcal{K}} y_j = Q^1. \tag{10}$$

The aggregate additional capacity in the buildout is Q^1 and is split among the operators of capacity: the hyperscalers and the neoclouds.

Compute requires a combination of real assets (real estate and power infrastructure) and IT hardware (GPU server racks, fiber optic cables, storage, and other IT peripherals). We will refer to the former as real estate and the latter as IT. Operators, both hyperscalers and neoclouds, must own all IT associated with the capacity they operate. Hyperscalers own some of the real estate and lease the rest from landlords. Let ϕ_j denote the fraction of owned real estate of operator j. Neoclouds do not own real estate: $\phi_k = 0$ for $k \in \mathcal{K}$. Landlords only own real estate assets, not IT. Let I_j be the investment in new IT and owned real estate by firm j. Let I_{L} be the landlords' investment in real estate:

$$I_j = \left[p_{IT}(Q^1) + b_{RE}\phi_j\right] y_j, \qquad I_{\mathsf{L}} = b_{RE} \sum_{j \in \mathcal{H} \cup \mathcal{K}} (1 - \phi_j) y_j, \tag{11}$$

where $p_{IT}(Q^1)$ is the price per unit of IT hardware and b_{RE} the price per unit of real estate.

3.5 Financing the Buildout

Firms take on new obligations associated with the buildout. For counterparties $i, j \in \mathcal{I}$, denote by $W_{ij}^{1,J}$ and $W_{ij}^{1,S}$ the face-value amounts of the new junior and senior obligations associated with the buildout between the firms inside the modeled network.

For the lab firms, the new junior obligations constitute service contracts for compute with the operators. Let p_{ja} denote the promised compute-service price per GW charged by operator j to lab a, expressed in date-zero-equivalent units. The face value of the lab's obligations for its chosen new compute quantity q_a^1 is:

$$W_{\cdot a}^{1,J} = R q_a^1 \sum_{j \in \mathcal{H} \cup \mathcal{K}} \omega_{ja} p_{ja}.$$

The hyperscalers buy some compute from the neoclouds as part of the buildout. Let p_{kh} denote the promised compute-service price per GW charged by by neocloud k to hyperscaler h, in date-zero-equivalent units. The face value of the hyperscaler h's obligation to neocloud k is:

$$W_{kh}^{1,J} = R p_{kh} \beta_{kh} d_h.$$

Operators also take on new rental commitments to the data center landlord sector L. Let $r_{\mathsf{L}i}$ denote the promised rental charge per GW paid by operator i, and let $x_{\mathsf{L}i}^1$ denote the additional capacity it leases. The new lease obligation is:

$$W_{\mathsf{L}i}^{1,J} = R r_{\mathsf{L}i} x_{\mathsf{L}i}^1, \qquad i \in \mathcal{H} \cup \mathcal{K}.$$

To pay for IT hardware purchases and/or real estate associated with the buildout, hyperscalers, neoclouds, and landlords use cash, newly-issued debt, and newly-issued equity. Let $D_i^{1,\text{out},S}$ denote the face value of new secured debt owed by firm i to outside lenders. This debt is subject to a loan-to-cost (LTC) limit:

$$D_i^{1,\text{out},S} \leq \lambda_i I_i, \qquad i \in \mathcal{H} \cup \mathcal{K} \cup \mathcal{L}. \tag{12}$$

In the baseline model, there is a single borrowing limit λ_i that is the same across firms and across collateral (real estate and IT). Crucially, this new financial debt to outside lenders is non-recourse to legacy assets; it is secured only by the cash flows attributable to the newly financed assets.[2]

A possible extension allows the LTC limit to differ for real estate and IT collateral. In that case, let I_i^{IT} and I_i^{RE} denote expenditure on new IT and owned real estate, and let λ_{IT}

[2]An alternative formulation would introduce a separate IT rental market in which external GPU owners borrow to purchase GPUs and lease them to neoclouds and hyperscalers. The current setup effectively consolidates the balance sheets of the neoclouds/hyperscalers and these GPU-owning financing vehicles, including GPU-backed securitization vehicles. The GPUs and the debt used to finance them therefore appear on the neocloud's modeled balance sheet, while rental payments between the neocloud and the vehicle cancel upon consolidation. The debt remains secured only by cash flows attributable to the newly financed assets.

and λ_{RE} denote their respective face-value LTC limits. The borrowing constraint in the alternative model is:

$$D_i^{1,\text{out},S} \leq \lambda_{IT} I_i^{IT} + \lambda_{RE} I_i^{RE}, \qquad I_i = I_i^{IT} + I_i^{RE}.$$

Neoclouds have $I_i^{RE} = 0$ and landlords have $I_i^{IT} = 0$. The effective λ_i for an integrated operator is the investment-cost-weighted average of the two limits. Differing LTC limits for real estate and IT could reflect the assets' different risk profiles or economic lives.

The actual loan proceeds L_i from new secured borrowing are endogenous as lenders' advances depend on the riskiness of the loan. We solve for L_i below.

To finance the buildout, hyperscalers, neoclouds, and landlords may also issue additional equity, subject to the constraint:

$$e_i = \max\{0, I_i - C_i^0 - L_i\} \leq \bar{e}_i, \qquad i \in \mathcal{H} \cup \mathcal{K} \cup \mathcal{L}. \tag{13}$$

where e_i are the proceeds from new equity issuance and $\bar{e}_i$ the maximum permitted new equity issuance.

3.6 The Balance Sheet Including the Buildout

Datacenter assets, including real estate and IT, are the sum of legacy assets and the new investment amount:

$$K_i = K_i^0 + I_i$$

for hyperscalers, neoclouds, and landlords, where it is understood that the neoclouds own only IT assets and landlords only real estate.

Post-buildout outside assets and receivables from counterparties outside the modeled network are H_i.

Post-buildout senior and junior obligations within the network are given by:

$$\begin{aligned} W_{ij}^J &= W_{ij}^{0,J} + W_{ij}^{1,J}, \\ W_{ij}^S &= W_{ij}^{0,S} + W_{ij}^{1,S} \end{aligned} \tag{14}$$

Post-buildout cash is:

$$C_i = C_i^0 + L_i + e_i - I_i \geq 0.$$

The law of motion for outside debt is:

$$D_i^{\text{out},S} = D_i^{0,\text{out},S} + D_i^{1,\text{out},S}.$$

Since all new debt is secured debt, the outside junior liabilities remain the same as legacy junior liabilities:

$$D_i^{\text{out},J} = D_i^{0,\text{out},J}.$$

Post-buildout equity in book-value terms is:

$$E_i^{BV} = E_i^{0,BV} + e_i$$

The post-buildout balance sheet, in book value terms, reads:

$$\begin{aligned} A_i^{BV} &= C_i + K_i + H_i + W_{i\cdot}^J + W_{i\cdot}^S + \sum_j O_{ij} E_j^{BV}, \\ D_i^{BV} &= D_i^{\text{out},J} + D_i^{\text{out},S} + W_{\cdot i}^J + W_{\cdot i}^S, \\ E_i^{BV} &= A_i^{BV} - D_i^{BV}. \end{aligned} \tag{15}$$

3.7 Risk Pricing: Market Values of Assets and Liabilities

New compute-service contracts, rental commitments, and debt obligations specify promised payments. Actual payments may fall short when obligors default. The pricing of these contracts and the loan proceeds reflect the present discounted, risk-adjusted values of the payments creditors receive. Likewise, the market value of equity is derived as the present discounted, risk-adjusted value of the residual cash flows accruing to shareholders. All values in this subsection are measured at date 0, after the buildout but before the state is realized, and are conditional on the entire vector of buildout choices. The existence of legacy obligations also enters the risk pricing, since ordinary legacy and buildout-associated obligations are pari-passu.

Market Value of Liabilities Let $\rho_j^J(s)$ denote the fraction paid on junior claims against firm j in state s, assuming equal proportional recovery within that priority class. Let ζ_i denote the risk-neutral expected repayment fraction on junior obligations of customer i:

$$\zeta_i = \sum_s \tilde{\pi}_s \rho_i^J(s). \tag{16}$$

Let $\rho_{ij}^S(s)$ denote the recovery fraction on the senior claim held by i against j. Different senior creditors can therefore receive different recovery fractions from the same borrower if the underlying collateral is different. Both fractions lie between zero and one. Let $W_{ij}^{MV,J}$

and $W_{ij}^{MV,S}$ denote the date-0 market values of the corresponding total claims:

$$\begin{aligned} W_{ij}^{MV,J} &= \sum_s \pi_s M_s W_{ij}^J \rho_j^J(s) = \left(\frac{1}{R}\sum_s \tilde{\pi}_s \rho_j^J(s)\right) W_{ij}^J = \frac{1}{R}\zeta_j W_{ij}^J, \\ W_{ij}^{MV,S} &= \sum_s \pi_s M_s W_{ij}^S \rho_{ij}^S(s) = \left(\frac{1}{R}\sum_s \tilde{\pi}_s \rho_{ij}^S(s)\right) W_{ij}^S. \end{aligned} \tag{17}$$

In similar fashion, we can mark-to-market the firm's junior and senior legacy obligations to outside parties. Let $\rho_i^{0,\text{out},S}(s)$ denote the recovery fraction on its inherited outside senior debt. The market value of legacy outside liabilities is:

$$D_i^{0,\text{out},MV} = \frac{1}{R}\sum_s \tilde{\pi}_s \left[D_i^{0,\text{out},J}\rho_i^J(s) + D_i^{0,\text{out},S}\rho_i^{0,\text{out},S}(s)\right].$$

Let $\rho_i^{1,\text{out},S}(s)$ denote the recovery fraction on firm i's new outside secured debt in state s. Fair pricing implies that loan proceeds equal the discounted, risk-adjusted value of repayments:

$$L_i = \left(\frac{1}{R}\sum_s \tilde{\pi}_s \rho_i^{1,\text{out},S}(s)\right) D_i^{1,\text{out},S}. \tag{18}$$

The debt is non-recourse: lenders recover only from the specified cash flows generated by newly financed assets, with no claim on legacy cash or outside-business income.

The recovery fractions $\rho_j^J(s)$, $\rho_{ij}^S(s)$, and $\rho_i^{1,\text{out},S}(s)$ are determined by the priority-of-payment rules, spelled out below.

The market value of outside liabilities includes the market value of legacy outside liabilities:

$$D_i^{out,MV} = D_i^{0,out,MV} + L_i.$$

Let $E_i(s)$ denote the date-1 residual cash flow paid to all shareholders of issuer i, after all senior and junior claimants have been paid according to the applicable priority and collateral rules. Let E_i^{MV} denote the firm's post-issuance market equity value. Let δ_j denote the fraction of equity in firm j sold to new outside investors and O_{ij} the diluted (post-issuance) ownership fraction in firm j held by firm i:

$$E_i^{MV} = \frac{1}{R}\sum_s \tilde{\pi}_s E_i(s), \qquad \delta_i = \frac{e_i}{E_i^{MV}}, \qquad O_{ij} = O_{ij}^0(1-\delta_j). \tag{19}$$

Issuing new equity requires $0 \le e_i \le E_i^{MV}$. Set $\delta_i = 0$ when $e_i = 0$; positive issuance requires $E_i^{MV} > 0$. The market value of the equity investment by firm i in issuer j is $O_{ij}E_j^{MV}$.

Market Value of Assets The market value H_i^{MV} includes outside financial assets, outside receivables, and the value of other operating cash flows. Let $X_i^H(s)$ denote the date-1 payoff of the outside assets and receivables in state s. These cash flows constitute legacy operating cash flows and the productivity-dependent benefits of additional compute use:

$$X_i^H(s) = X_i^{H,0}(s) + R\mathcal{V}^P(q_i^1; Q_{-i}^1, s).$$

Then the pricing dictates:

$$H_i^{MV} = \frac{1}{R}\sum_s \tilde{\pi}_s X_i^H(s) = \frac{1}{R}\sum_s \tilde{\pi}_s X_i^{H,0}(s) + \sum_s \tilde{\pi}_s \mathcal{V}^P(q_i^1; Q_{-i}^1, s) = H_i^{0,MV} + \sum_s \tilde{\pi}_s \mathcal{V}^P(q_i^1; Q_{-i}^1, s), \tag{20}$$

As explained below, the calibration will back out $H_i^{0,MV}$ from the observed stock market value of the firm. We parameterize the distribution of legacy outside cash flows across AI states using the function $g_i(s)$:

$$X_i^{H,0}(s) = RH_i^{0,MV} g_i(s), \qquad g_i(s) = \frac{1-\alpha_i+\alpha_i\theta_s}{\sum_t \tilde{\pi}_t(1-\alpha_i+\alpha_i\theta_t)}, \qquad \sum_s \tilde{\pi}_s g_i(s) = 1. \tag{21}$$

The parameter α_i governs exposure to the AI productivity state. A firm with $\alpha_i = 0$ has legacy financial assets whose time-1 cash flows do not depend on the state of AI.

We model data center assets (real estate and IT) similarly. Let $X_i^K(s)$ denote the residual payoff of the (legacy and newly built) data center assets in period 1. Then the market value of these assets is:

$$K_i^{MV} = \frac{1}{R}\sum_s \tilde{\pi}_s X_i^K(s)$$

We set this residual payoff $X_i^K(s) = 0, \forall s$, which implies $K_i^{MV} = 0$. We include the operating benefits of compute use in H, rather than K, because they accrue to the firm using compute, regardless of who owns the underlying hardware and facilities. In particular, AI labs earn these benefits despite owning no data-center assets. Operators' compute-service revenues and landlords' rental revenues are recorded separately as contractual receipts in W. Thus, K captures only the residual value of physical assets beyond the operating and contractual cash flows already accounted for. Allowing positive salvage values is a possible extension.

If cash earns the risk-free rate R, its market value is simply C_i.

Balance Sheet in Market Values Let A_i^{MV} and D_i^{MV} denote total assets and liabilities in market-value terms, post-buildout:

$$\begin{aligned} A_i^{MV} &= C_i + H_i^{MV} + K_i^{MV} + W_{i\cdot}^{MV,J} + W_{i\cdot}^{MV,S} + \sum_{j\in\mathcal{I}} O_{ij} E_j^{MV}, \\ D_i^{MV} &= D_i^{out,MV} + W_{\cdot i}^{MV,J} + W_{\cdot i}^{MV,S}, \\ A_i^{MV} &= D_i^{MV} + E_i^{MV}. \end{aligned} \tag{22}$$

Backing Out Legacy Outside Assets In the calibration below, we pin down the market value of legacy outside assets, which includes the value from operational cash flows from business activities unrelated to AI, by matching the incumbent market value of equity in the model to the observed market value of equity in the data. In particular, we jointly determine the vector $\boldsymbol{H}^{0,MV}$ and the equilibrium buildout $\boldsymbol{q}^{1,*}$ to satisfy

$$E_i^{MV}\big(\boldsymbol{H}^{0,MV}, \boldsymbol{q}^{1,*}\big) - e_i\big(\boldsymbol{H}^{0,MV}, \boldsymbol{q}^{1,*}\big) = E_i^{\text{obs}}, \qquad i \in \mathcal{I}. \tag{23}$$

Here E_i^{obs} is observed incumbent equity value, using the stock market valuation on June 30, 2026 for the publicly-listed companies and the most recent private-market equity valuation for the non-listed companies. Given an empirical counterpart to all other assets and all liabilities in 22, this allows us to back out $H_i^{0,MV}$ for every firm i. The implicit assumption is that shareholders in the data anticipate the modeled buildout.

3.8 Decentralized Equilibrium

Generalized Nash Equilibrium Given legacy assets and liabilities, a generalized Nash equilibrium consists of a compute allocation $\boldsymbol{q}^{1,*}$ and associated hardware prices p_{IT}, compute-service prices $\{p_{ja}, p_{kh}\}$, and rents $\{r_{\mathsf{L},j}\}$. Each firm $i \in \mathcal{A} \cup \mathcal{H}$ chooses additional compute capacity to maximize incumbent shareholder value (6), holding other firms' investment choices fixed and anticipating the resulting prices, financing terms, and network payments. Hardware prices follow the margin schedule specified below, while new compute-service and rental contracts are priced competitively to cover suppliers' risk-adjusted costs.

No firm can increase incumbent shareholder value through a feasible unilateral change in its compute investment. At each potential deviation, we recompute hardware prices, individualized break-even contract prices, routed customer demand, capital expenditure, contractual obligations, risky loan proceeds, equity dilution, and network payments. Feasibility requires

that the resulting allocation satisfy the debt and equity financing constraints (12) and (13) throughout the network. Because a customer's expansion uses its suppliers' financing capacity, each firm's feasible investment set depends on other firms' choices.

Prices for IT Hardware Let b_{IT} denote the resource cost of producing IT hardware sufficient to support one GW of compute capacity. Hardware producers, $b \in \mathcal{B}$, have market power and charge a price $p_{IT}(Q^1)$ above production cost. We assume that their profit margin declines with aggregate new buildout Q^1, subject to a lower bound of zero:

$$\mu(Q^1) = \max\{0, \mu_H - b_\mu Q^1\}, \qquad p_{IT}(Q^1) = \frac{b_{IT}}{1 - \mu(Q^1)}.$$

This specification captures, in reduced form, the idea that high margins attract entry and expansion of competing chip production capacity. A larger buildout therefore reduces margins and hardware prices. For $Q^1 \geq \mu_H / b_\mu$, chips are sold at production cost.

Hardware producers in $\mathcal{B}$ earn the following aggregate profits:

$$\Pi_{IT}(Q^1) = \left[p_{IT}(Q^1) - b_{IT}\right] Q^1. \tag{24}$$

Producer b receives $\nu_b \Pi_{IT}(Q^1)$ in cash at date 0, where ν_b is its sales share and $\sum_{b \in \mathcal{B}} \nu_b = 1$. Operator and landlord capital expenditure satisfy:

$$\sum_{j \in \mathcal{H} \cup \mathcal{K}} I_j + I_{\mathsf{L}} = [p_{IT}(Q^1) + b_{RE}]Q^1 = C^S(Q^1) + \Pi_{IT}(Q^1), \tag{25}$$

where

$$C^S(Q^1) = (b_{RE} + b_{IT})Q^1 \tag{26}$$

is the social cost of building additional compute Q^1. Hardware profit $\Pi_{IT}(Q^1)$ is a transfer between firms in the AI ecosystem, rather than a real resource cost.

Competition in Rental Market We assume the market for new data center leases is competitive. A rental contract between an operator and a landlord is an example of a junior obligation. Applying (17) to the new lease gives:

$$W_{\mathsf{L}i}^{1,MV,J} = r_{\mathsf{L}i} x_{\mathsf{L}i}^1 \zeta_i. \tag{27}$$

The factor R in each date-1 face value cancels risk-free discounting. A promised charge in date-0 units must be multiplied by ζ_i to obtain its market value.

Under the assumption of no residual property value for the data center, discussed above, a competitive landlord requires the market value of a new rental contract to cover the construction cost:

$$r_{\mathsf{L}i}\zeta_i = b_{RE} \quad \Rightarrow \quad r_{\mathsf{L}i} = \frac{b_{RE}}{\zeta_i}. \tag{28}$$

Rents are set competitively, compensating the landlord for tenant credit risk. Riskier tenants pay higher rents in equilibrium.

Competition in Compute Market We assume the market for new compute contracts is competitive. For newly-negotiated compute contracts, the pricing rule in (17) gives

$$W_{ja}^{1,MV,J} = p_{ja}\omega_{ja}q_a^1\zeta_a, \tag{29}$$

Let κ_j^{dir} denote operator j's cost of directly supplying one GW. Let κ_j^{del} denote its cost of delivering the compute, including subcontracting.

$$\begin{aligned}
\kappa_j^{\text{dir}} &= p_{IT}(Q^1) + \phi_j b_{RE} + (1-\phi_j) r_{\mathsf{L}j}\zeta_j, \\
\kappa_h^{\text{del}} &= \left(1 - \sum_k \beta_{kh}\right) \kappa_h^{\text{dir}} + \sum_k \beta_{kh} p_{kh}\zeta_h.
\end{aligned} \tag{30}$$

For neoclouds $k \in \mathcal{K}$, $\kappa_k^{\text{del}} = \kappa_k^{\text{dir}}$. These expressions use the market value of the operator's own rental and subcontracting payments, rather than the promised amounts. Competitive compute contracts satisfy:

$$p_{kh}\zeta_h = \kappa_k^{\text{dir}}, \qquad p_{ja}\zeta_a = \kappa_j^{\text{del}}. \tag{31}$$

Substituting in for the competitive rental rate from (28) and subcontracting break-even conditions allows us to solve for the competitive compute contract and subcontract prices:

$$p_{ja} = \frac{p_{IT}(Q^1) + b_{RE}}{\zeta_a}, \qquad p_{kh} = \frac{p_{IT}(Q^1) + b_{RE}}{\zeta_h}. \tag{32}$$

Riskier customers end up paying a higher price per unit of compute to compensate the operator for the risk of non-payment. Hence, promised compute prices differ across customers for a given provider. The operator's own credit risk raises its promised rental and subcontracting payments, but competition ensures that the market value of those payments equals the underlying supply costs. The compute price charged to the lab therefore depends only on the lab's credit risk, which determines how much it must promise to cover those costs.

3.9 Optimal Buildout

The social benefit of having Q^1 units of compute in operation at time 1 in state s generates output $\mathcal{V}^S(Q^1; s)$:

$$\mathcal{V}^S(Q^1; s) = \theta_s \left(a_V Q^1 - \frac{b_V}{2}(Q^1)^2 \right), \tag{33}$$

with marginal social benefit:

$$\frac{\partial \mathcal{V}^S(Q^1; s)}{\partial Q^1} = \theta_s(a_V - b_V Q^1),$$

where a_V denotes the marginal benefit of the first increment of compute and $b_V > 0$ the rate at which the marginal benefit falls with aggregate compute. There are decreasing returns to compute in society.

The social cost of building Q units of compute is given by (26). The marginal social cost of compute is constant.

The social planner expands compute capacity until the expected marginal social benefit equals the marginal social cost of compute:

$$\tilde{\theta}(a_V - b_V Q^{1,*}) = b_{RE} + b_{IT},$$

where $\tilde{\theta} = \sum_s \tilde{\pi}_s \theta_s$; the social planner uses the same risk-neutral probabilities as the investors in the model.

3.10 Buildout at the Decentralized Equilibrium

Four mechanisms can drive a wedge between the decentralized buildout and the planner's allocation. The first two concern operating benefits and costs; the remaining two arise from financing. The operating wedges described below isolate firms' direct investment incentives. Cross-holdings and contractual links can modify these incentives in the full shareholder-value problem.

First, hardware prices exceed resource costs whenever chip producers earn a positive profit margin. The planner counts the resource cost of hardware production, but treats hardware profits as transfers within the economy. For a firm choosing q_i^1, with $Q^1 = q_i^1 + Q_{-i}^1$, the private supply cost is $C_i^P = q_i^1[p_{IT}(Q^1) + b_{RE}]$. Away from the kink where the profit margin reaches zero, its marginal cost is

$$\frac{\partial C_i^P}{\partial q_i^1} = b_{RE} + \frac{b_{IT}}{1 - \mu(Q^1)} + \frac{q_i^1 b_{IT} \mu'(Q^1)}{[1 - \mu(Q^1)]^2}.$$

The hardware price premium discourages investment, holding prices fixed. Before the margin floor binds, however, $\mu'(Q^1) < 0$: expansion lowers the price of hardware for the firm's own buildout, partially offsetting this premium. It also lowers other builders' hardware expenditure, a benefit the investing firm does not internalize. This benefit is a pecuniary externality: lower hardware prices redistribute surplus from chip producers to buyers rather than reduce aggregate resource costs.

Second, a firm's expansion reduces the operating benefits earned by other compute users. The firm accounts for decreasing returns to its own investment, but does not generally internalize the losses imposed on other users. Differentiating (7) gives

$$\frac{\partial \mathcal{V}^P(q_i^1; Q_{-i}^1, s)}{\partial q_i^1} = \theta_s \left(a_V - b_V q_i^1 - \frac{b_V}{2} Q_{-i}^1 \right).$$

At the same aggregate allocation, the private marginal benefit therefore exceeds the social marginal benefit by

$$\frac{\partial \mathcal{V}^P}{\partial q_i^1} - \frac{\partial \mathcal{V}^S}{\partial Q^1} = \frac{\theta_s b_V}{2} Q_{-i}^1.$$

This wedge creates an incentive to overinvest relative to the planner, all else equal.

Third, firms maximize incumbent shareholder value, which can create incentives to shift risk onto inherited creditors. With limited liability, shareholders receive the residual gains in favorable states, while creditors can bear part of the losses in unfavorable states. An investment that reduces the value of inherited claims can therefore benefit shareholders even when it reduces total firm value. Competitive pricing of new debt compensates new lenders for the risk they assume, but does not compensate inherited creditors whose contractual terms are fixed. This channel can encourage excessive or excessively risky investment.

Fourth, financing constraints can prevent otherwise desirable investment. A feasible buildout requires service and rental prices that allow suppliers to break even, loan proceeds consistent with collateral recoveries, and equity issuance consistent with investor participation and issuance limits. These conditions must hold jointly throughout the network: a customer's expansion may require investment and financing by its suppliers. Loan-to-cost limits and equity issuance caps can therefore constrain aggregate buildout below the planner's allocation.

Because these mechanisms can push investment in different directions, whether the decentralized equilibrium exhibits overinvestment or underinvestment relative to the planner is a quantitative question. Section 5.2 reports the numerical equilibrium allocation and prices.

3.11 Shock Realizations at Time 1

In each state at time 1, firms' cash flows and contractual receipts determine creditor recoveries and residual equity through the network of claims. Firm i's total resources in state s at time one equal

$$A_i(s) = RC_i + X_i^H(s) + X_i^K(s) + \sum_{j\in\mathcal{I}} \left[W_{ij}^J \rho_j^J(s) + W_{ij}^S \rho_{ij}^S(s)\right] + \sum_{j\in\mathcal{I}} O_{ij} E_j(s). \tag{34}$$

Let $\chi_h = 1 - \sum_k \beta_{kh}$ be the fraction of hyperscaler h's demand that it operates directly. The cash flows pledged to new outside lenders are:

$$\begin{aligned}
\Gamma_h(s) &= \chi_h \left[R\mathcal{V}^P(q_h^1; Q_{-h}^1, s) + \sum_a W_{ha}^{1,J} \rho_a^J(s)\right]_+, \\
\Gamma_k(s) &= \sum_a W_{ka}^{1,J} \rho_a^J(s) + \sum_h W_{kh}^{1,J} \rho_h^J(s), \\
\Gamma_{\mathsf{L}}(s) &= \sum_{j\in\mathcal{H}\cup\mathcal{K}} W_{\mathsf{L}j}^{1,J} \rho_j^J(s).
\end{aligned} \tag{35}$$

Hyperscalers pledge income attributable to their directly operated new IT and owned facilities, neoclouds pledge new service receipts generated by their IT, and landlords pledge rents on newly financed facilities.

New lenders receive the minimum of the promised face value and the collateral value:

$$D_i^{1,\text{out},S} \rho_i^{1,\text{out},S}(s) = \min\{D_i^{1,\text{out},S}, \Gamma_i(s)\}.$$

Unpaid new non-recourse debt has no claim on other resources and does not become a junior liability. Let $S_i(s)$ denote total payments to the secured creditors of firm i in state s. After payments to new non-recourse lenders, legacy senior claims retain priority over remaining corporate resources, with proportional recovery within the legacy senior class. Let D_i^J denote total junior liabilities, including those owed to outside creditors. Remaining resources are distributed proportionally to junior creditors. Shareholders receive the residual:

$$\begin{aligned}
J_i(s) &= \min\{D_i^J, A_i(s) - S_i(s)\}, \\
\rho_i^J(s) &= \frac{J_i(s)}{D_i^J}, \qquad E_i(s) = [A_i(s) - S_i(s) - D_i^J]_+.
\end{aligned} \tag{36}$$

Here $[z]_+ = \max\{z, 0\}$, and $\rho_i^J(s) = 1$ when $D_i^J = 0$. Feasibility requires $0 \le S_i(s) \le A_i(s)$. Payments, recoveries and equity payoffs are determined jointly across firms because each firm's resources depend on payments received from others.

4 Data and Calibration

4.1 Data

An important contribution of this paper is to measure the obligations that the most important firms in the AI buildout have vis-a-vis one another and vis-a-vis the firms and investors outside the network. The complexity and opacity of the financing arrangements makes this a difficult but important task.

We assemble a firm-level data set of the financial and contractual linkages underlying the AI buildout. The data set combines balance sheets, contractual commitments, equity ownership, and data-center capacity to measure exposures both within the modeled network and to outside creditors, suppliers, and investors. The financial reference date (date zero) is June 30, 2026. Appendix A reports the sources, construction, and assumptions for each input. Dollar amounts are expressed in billions of US dollars and capacity in GW of facility power.

Coverage and sources. The network comprises eleven named firms and an aggregate data-center landlord sector. The AI labs $\mathcal{A}$ are OpenAI and Anthropic; the hyperscalers $\mathcal{H}$ are Microsoft, Amazon, Alphabet, Meta, Oracle, and the consolidated SpaceX/xAI/X group; the neoclouds $\mathcal{K}$ are CoreWeave and Nebius; and the hardware producer set $\mathcal{B}$ consists of NVIDIA. The landlord set $\mathcal{L} = \{\mathsf{L}\}$ contains a single aggregate sector representing third-party owners of data-center facilities. We plan to expand the set of named entities in future iterations.

Our primary sources are SEC filings, including annual and quarterly reports, registration statements, ownership disclosures, and contractual exhibits. We supplement these with company announcements, earnings calls, and financial press reports. The news articles are a particularly valuable additional source of information for private firms and contracts whose counterparties or amounts are not disclosed in financial statements.

Market capitalizations and balance sheet items come from CRSP and Compustat. Pitchbook provides estimated market values for private companies such as OpenAI and Anthropic as well as for strategic cross-investments in AI companies. Cleanview provides project-level data-center capacity. Epoch AI provides estimates of chip ownership and additions.

We use the latest balance sheet dated on or before the reference date. Later disclosures are used when they describe positions or contracts existing at that date.

Contractual and ownership linkages. We record compute-service commitments between AI labs and cloud providers, subcontracting commitments between hyperscalers and neoclouds, hardware-related trade credit, and rental obligations to the landlord sector. The inherited ledger includes booked liabilities and signed but unfulfilled commitments, including uncommenced leases and future purchase, construction, and compute payments. These uncommenced obligations are part of the legacy obligations rather than the buildout since they are associated with data-center projects that are already operational or under construction.

Each identified bilateral claim is carved out of the obligor's total obligations; the remaining amount is assigned to counterparties outside the network. This procedure avoids counting the same obligation twice. We infer outside residuals only when aggregate disclosures and named contracts have consistent coverage. Equity cross-holdings are recorded separately using disclosed ownership or common-equivalent interests. Where amounts or counterparties are unavailable, the appendix identifies the assumptions used.

NVIDIA's hardware receivables require an additional allocation. Because its direct customers often are equipment manufacturers and intermediaries, we attribute receivables to economic end users using NVIDIA's receivables-to-revenue ratio, customer segment disclosures, Epoch AI chip additions, and equipment additions reported by neoclouds.

Legacy balance sheets and valuation. We collect cash, financial debt, lease liabilities, trade payables, and signed commitments, distinguishing secured debt from ordinary obligations according to the model's priority structure. Undiscounted commitments are converted to present-value proxies using a 5% annual discount rate, a fifteen-year horizon for leases, and a five-year horizon for purchase, construction, and compute commitments. Current lease liabilities reported on balance sheets are already discounted to present value and are therefore used without further discounting.

Public-company equity values use reference-date prices and relevant share counts; private-company targets use disclosed financing valuations and common-equivalent shares.

Physical capacity and the landlord sector. We measure owned facility capacity from Cleanview projects operating or under construction in its September 12, 2026 extract. US capacity is scaled to worldwide capacity using a common scaling factor of 1.2. Leased capacity is inferred from rental obligations, allowing us to construct each operator's owned-site share. Because the landlord sector has no available consolidated financial statements,

its property value, debt, and equity are constructed from capacity and explicit assumptions about development costs, rents, operating expenses, and leverage.

4.2 Calibration

Table 1 lists the main parameters. The model's time discount rate is $R = 1.05$.

The AI productivity grid is symmetric around 1 with one very good state (1.95) and one very bad state (0.05). Each of these two states occurs with 25% probability. Mean productivity is 1.

Table 1: Parameter values.

Notation	Interpretation	Calibration
R	Gross risk-free model-period return	1.05
θ_s	Productivity grid	$(1.95, 1, .05)$
π_s	Physical probabilities	$(.25, .50, .25)$
$\tilde{\pi}_s$	Risk-neutral probabilities	$(.15, .50, .35)$
a_V	Initial marginal operating benefit	\$80B/GW
b_V	Decline in marginal benefit per GW	\$.30B/GW^2
b_{RE}	Real estate and power infrastructure cost	\$13.33B/GW
b_{IT}	Real IT hardware production cost	\$8B/GW
μ_H	Initial hardware profit margin	.70; lower bound is zero
b_μ	Margin reduction per added GW	.002 per GW
λ_j	New secured loan face/capex	.80 for all operators
$\bar{e}_j$	New operator-equity allowance	\$275 billion per operator; landlord uncapped
ϕ_j	Share of RE owned	Fixed proxy; 0 for neoclouds

The pricing measure shifts 10% probability mass from the highest to the lowest state, consistent with risk-averse investors. Consequently, $\tilde{\theta} = .81 < \bar{\theta} = 1$. The SDF is largest in the low state and satisfies $\sum_s \pi_s M_s = 1/R$.[3]

[3]These risk prices are illustrative assumptions rather than estimates from security returns. They could be chosen to match evidence from CDS spreads in a future draft.

Turning to the technology parameters, at the middle productivity state (also equal to average productivity), the marginal social benefit of compute is $80 - .30Q$. It is \$80B per GW for the first increment of compute, falls to \$50B/GW at a 100 GW buildout, and to \$20B/GW at 200 GW. The marginal benefit becomes negative above 267 GW.[4] The total operating benefit becomes negative above 533 GW.

Building datacenters costs \$13.33B/GW in real estate and power infrastructure costs (Van Nieuwerburgh, 2026). The IT hardware producers' profit margin is 70% for the first unit of compute and falls to 30% at a 200 GW buildout. It reaches zero at 350 GW. This makes the initial IT purchase price p_{IT} equal to \$26.67B/GW. Adding the \$13.33B/GW in real estate costs gives total capital expenditure of \$40 billion per GW. This is the cost estimate used by Van Nieuwerburgh (2026). The price of IT hardware falls to the social resource cost of \$8B/GW for a large buildout, at which point the total private construction cost becomes \$21.33B/GW. The social cost is \$21.33B/GW at every quantity.

In the benchmark calibration, the maximum loan face value-to-cost ratio is λ_i =80%. Loan proceeds equal the risk-adjusted value of pledged repayments, so an 80% face-to-cost ratio implies that lenders advance less than 80% of capital expenditures. We assume that all borrowers, including the landlord sector, issue new secured debt at the full face-value LTC limit. The extension with separate IT and real-estate borrowing limits could, for example, use LTC ratios of 60% and 80%, respectively.

Each operator may issue up to $\bar{e}_j$ = \$275B of new equity; the landlord sector has no equity-issuance cap.

AI lab procurement shares ω_{ja} and compute subcontracting shares β_{kh} in the buildout are set at their observed legacy values, which are documented in Appendix Tables A9 and A10, respectively. This calculation uses data from both current and future commitments.

Owned-site fractions ϕ_j in the buildout are set at their observed legacy values. The latter are measured from owned datacenter facilities that are operational or under construction and where the developer is the named developer, and from leased capacity, which is inferred from reported rental obligations. The details are in Appendix A.5.

The AI exposure of legacy outside assets for firm i, α_i, is reported in Appendix Table A11.

[4]Negative marginal social benefits could arise from pollution, noise, water scarcity, etc. arising from a large buildout.

4.3 The Legacy Network

The inherited network is concentrated around several large customer relationships. The last two columns of Table 2 report total legacy obligations D_i^0 and their financial-debt component $D_i^{0,\text{fin}}$, including senior and junior debt, as of June 30, 2026. OpenAI has \$666.0 billion of ordinary obligations, including \$659.3 billion of compute-service commitments; Microsoft and Oracle together hold approximately three quarters of these compute-related claims. Amazon and Alphabet account for 85.3% of Anthropic's named compute-service commitments; Anthropic's total ordinary obligations amount to \$454.7 billion. The top panel of Figure 1 depicts these lab procurement relationships. Payment shortfalls by the labs can reduce their suppliers' resources and, in turn, their ability to honor their own obligations.

The second panel shows compute subcontracting relationships between hyperscalers and neoclouds. CoreWeave has \$18.2 billion of secured debt and \$61.6 billion of ordinary obligations, while Nebius has \$24.2 billion of ordinary obligations. The network also links operators to hardware producers and property owners. NVIDIA's eight modeled receivables total \$28.34 billion. The landlord sector holds rental claims across operators and has \$632.95 billion of secured debt and \$188.28 billion of equity. The appendix reports the complete bilateral ledger.

In aggregate, the modeled entities have \$5.22 trillion of legacy obligations, including \$1.21 trillion of financial debt, against \$20.55 trillion of reference equity value. These are sums across entities, before eliminating internal claims and cross-held equity. Obligations are distributed unevenly relative to equity values. Total obligations amount to approximately 2.5% of equity for SpaceX/xAI/X and 3.8% for NVIDIA, and 21–23% for Alphabet, Microsoft, and Amazon. The ratios are higher for Meta (44%), Anthropic (61%), OpenAI (78%), Oracle (92%), and CoreWeave (145%). These comparisons include operating commitments as well as financial debt and therefore measure a broader exposure than conventional debt leverage.

Although aggregate equity value substantially exceeds recorded obligations, this comparison alone does not establish resilience: equity valuations depend on future earnings, and payment shortfalls can propagate through the network. The concentration of obligations motivates the firm-level stress tests below, first for the legacy network and then incorporating the additional buildout.

Panel (c) of Figure 1 shows cross-equity holdings, which create an additional channel of transmission: a decline in one firm's equity value reduces the asset values of other modeled firms that hold stakes in it.

Firm	$H_i^{0,MV}$	E_i^{obs}	$E_i^{MV} - e_i$	E_i^{MV}	e_i	D_i^0	$D_i^{0,\text{fin}}$
OpenAI	639.00	852.00	852.00	852.00	0.00	666.05	0.00
Anthropic	607.79	739.90	739.90	739.90	0.00	454.68	0.00
Microsoft	2765.27	2771.00	2771.00	3042.93	271.93	596.93	40.29
Amazon	2621.53	2564.00	2564.00	2834.96	270.96	597.01	132.22
Google	4782.42	4348.30	4348.30	4606.84	258.54	906.27	100.16
Meta	1717.27	1435.30	1435.30	1701.02	265.72	633.20	83.66
Oracle	31.86	422.10	422.10	690.89	268.79	386.31	129.54
CoreWeave	56.09	54.90	54.90	63.49	8.59	79.81	35.07
Nebius	69.08	80.90	80.90	80.90	0.00	24.18	8.55
NVIDIA	4571.69	4846.40	4846.40	4846.40	0.00	182.55	8.47
SpaceX / xAI / X	1940.32	2251.30	2251.30	2501.51	250.21	55.33	38.43
Landlord	-81.91	188.28	188.28	852.77	664.50	632.95	632.95
Total	19720.40	20554.38	20554.38	22813.61	2259.23	5215.27	1209.36

Table 2: Equity matching and legacy obligations, in billions of dollars. Outside-income levels are jointly fitted with investment, retaining the stated AI exposures. Incumbent equity is post-issuance equity less newly issued equity; the maximum matching error is below 10^{-4} billion dollars. D_i^0 denotes total legacy senior and junior obligations to counterparties inside and outside the network, at date-0 present values, including financial debt, payables, commenced and uncommenced leases, and signed purchase, construction and compute commitments. $D_i^{0,\text{fin}}$ denotes legacy unsecured and secured financial debt, which is included in D_i^0. Sources and construction are documented in Appendix A. Total legacy obligations comprise approximately \$2.00 trillion owed within the network and \$3.22 trillion owed outside it. Totals use underlying unrounded amounts.

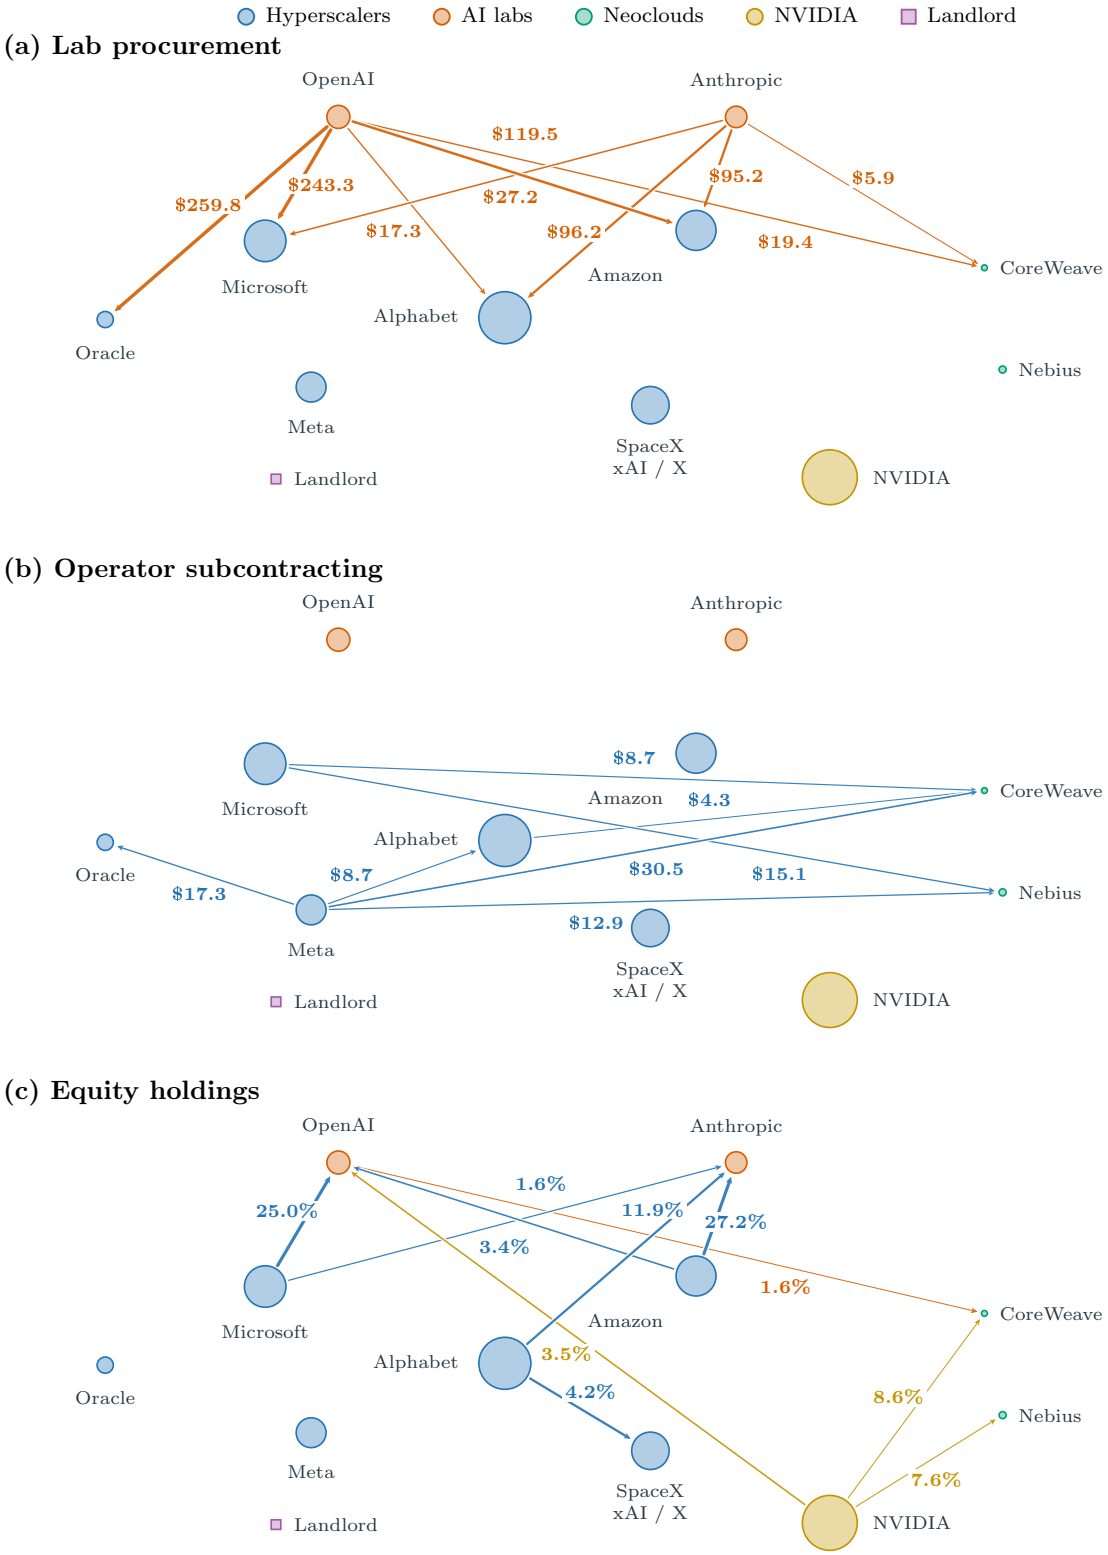


Figure 1: The legacy network on June 30, 2026.

Notes. The graph displays three types of networks among the key companies involved in the AI buildout: compute procurement contracts from the AI labs (panel a), subcontracting of compute (panel b), and strategic equity investments (panel c). See Appendix A.3for details.

The ability to honor these obligations depends on the resources available when AI productivity falls. We infer legacy outside-income values $H_i^{0,MV}$ jointly with equilibrium investment to match observed incumbent equity values. These residual valuations capture cash flows from existing businesses and outside assets, excluding the operating benefits of additional compute. For diversified hyperscalers, they include substantial earnings from activities such as advertising, social media, and e-commerce. Their proximity to observed equity values for some firms does not imply that other balance-sheet components are small: large contractual assets, investment benefits, and liabilities can substantially offset one another.[5]

Matching equity values determines the level of legacy outside income conditional on the assumed exposures α_i, which govern its distribution across AI states. Under these assumptions, Microsoft's legacy outside-income payoff falls from \$3.11 trillion in the middle state to \$2.08 trillion in the low state, while OpenAI's falls from \$828 billion to \$41 billion. Table A12 reports the complete state-contingent payoffs. These differences shape firms' ability to absorb losses on their network claims: outside income that remains resilient can support payments when customers default, whereas income that falls alongside customer receipts provides a weaker buffer. The stress tests below assess these sources of vulnerability jointly with contractual obligations, collateral arrangements, and cross-equity holdings.

5 Results

Having described the legacy exposures, we now turn to the model's predictions for the buildout. We contrast the optimal to the equilibrium buildout.

5.1 Unconstrained Social Planner

The socially optimal buildout equates the risk-adjusted marginal benefit of additional compute, which declines with aggregate capacity, to the constant marginal resource cost of hardware and data-center infrastructure. The planner's first-order condition from Section 3.9 gives:

$$\tilde{\theta}(a_V - b_V Q^*) = b_{IT} + b_{RE}, \qquad Q^* = \frac{80 - (8 + 13.33)/.81}{.30} = 178.88 \text{ GW}. \tag{37}$$

[5]Microsoft's inferred legacy outside-asset value is close to its observed equity value because large gross positions nearly offset: new compute benefits, contractual receivables, and equity stakes total approximately \$1.72 trillion, against \$1.71 trillion in liabilities and newly issued equity.

At this allocation, risk-adjusted operating benefits are \$7.70 trillion and resource costs are \$3.82 trillion, leaving \$3.89 trillion of surplus.[6]

5.2 Decentralized Equilibrium

As discussed in Section 3.8, we solve for a generalized Nash equilibrium in investment choices, accounting for how each firm's investment affects individualized compute and rental prices, loan proceeds, and financing feasibility. We jointly calibrate legacy outside-income values so that the equilibrium matches all twelve reference equity targets. The resulting allocation admits no profitable feasible unilateral deviation on a grid with 2.5-GW increments from zero through 400 GW for each decision maker.

The benchmark equilibrium builds 277.5 GW, 55% above the planner's allocation, at a capital expenditure of \$6.30 trillion. Risk-adjusted resource surplus is \$2.71 trillion, or \$1.18 trillion below the social optimum. Thus, competitive contract pricing and the imposed financing limits do not prevent overinvestment in this calibration. The result is consistent with firms' incentives to shift risk onto inherited creditors and their failure to fully internalize the reduction in other users' operating benefits caused by additional compute capacity.

Who uses and builds compute. Table 3 shows how each firm participates in the equilibrium buildout, distinguishing between the use and the ownership of compute capacity. OpenAI and Anthropic use 67.5 and 40.0 GW, respectively, all supplied by operators. Oracle installs 76.36 GW, Microsoft 47.09 GW and Amazon 39.09 GW; each serves lab customers as well as its own activity. Meta uses 40.0 GW and operates 38.08 GW, subcontracting 4.79% of its demand. CoreWeave and Nebius together install 8.84 GW. Their capacity is induced by customer orders, with the subcontracting shares routing most hyperscaler demand to the hyperscalers' own installations. The landlord supplies 151.28 GW of facilities; over half (54%) of data center real estate needed for the buildout is rented rather than owned by the operators. Compute use, IT ownership and facility ownership are three views of the same buildout, as the equal column totals in Table 3 show.

[6] Using physical rather than risk-neutral probabilities ($\bar{\theta}$ instead of $\tilde{\theta}$) would raise the optimal buildout to 195.56 GW.

Table 3: Additional compute use, IT ownership and real-estate ownership in the benchmark buildout, in GW.

Firm	Compute use	IT owned	RE owned
OpenAI	67.50	0.00	0.00
Anthropic	40.00	0.00	0.00
Google	17.50	35.54	27.42
Meta	40.00	38.08	18.85
SpaceX / xAI / X	32.50	32.50	31.51
Amazon	10.00	39.09	31.09
Microsoft	20.00	47.09	17.14
Oracle	50.00	76.36	0.20
CoreWeave	0.00	6.33	0.00
Nebius	0.00	2.51	0.00
NVIDIA	0.00	0.00	0.00
Landlord	0.00	0.00	151.28
Total	277.50	277.50	277.50

IT owned includes capacity serving customers after subcontracting. RE owned is the compute capacity supported by newly owned buildings and power infrastructure. Operators own their IT even at leased sites; the landlord owns those facilities. All entries concern additional capacity, not legacy stocks.

Prices and financing. At the equilibrium buildout, NVIDIA's profit margin falls to 14.5% and the IT hardware purchase price $p_{IT}(Q^1)$ is \$9.36 billion per GW, giving a total cost of \$22.69 billion per GW after adding the cost of real estate. To cover this cost, compute-service providers require OpenAI and Anthropic to promise \$33.25 and \$32.64 billion per GW, respectively. Their 46.55% and 43.87% price premiums compensate suppliers for payment risk. Similarly, Oracle's promised rent carries a 51.31% credit premium. CoreWeave also pays a rent premium. Table 4 reports the corresponding repayment discounts and charges.

Table 4: Equilibrium individualized contract prices.

Customer	ζ_i	Compute	Rent	Credit premium (%)
OpenAI	0.682	33.25	–	46.55
Anthropic	0.695	32.64	–	43.87
Microsoft	1.000	22.69	13.33	0.00
Amazon	1.000	–	13.33	0.00
Google	1.000	22.69	13.33	0.00
Meta	1.000	22.69	13.33	0.00
Oracle	0.661	–	20.17	51.31
SpaceX / xAI / X	1.000	–	13.33	0.00
CoreWeave	0.800	–	16.66	24.93
Nebius	1.000	–	13.33	0.00

Compute and rent are promised charges in billions of date-zero-equivalent dollars per GW; contractual date-one faces multiply these amounts by R. Compute covers new lab orders and active wholesale purchases; a dash indicates no such purchase. The credit premium is $100(1/\zeta_i - 1)$ percent over risk-adjusted supply cost, not an expected profit margin. Risk-adjusted compute cost is the IT purchase price plus construction cost; expected rent equals construction cost.

New loan proceeds finance \$3.73 trillion of the expansion, new equity \$2.26 trillion and existing cash \$0.31 trillion. Nearly 60% of the buildout is financed with debt. Similar to compute prices and rents, debt prices reflect the credit risk of the issuer. Table 5 shows the loan proceeds per dollar of promised payment in the first column. It reflects both loan proceeds and time discounting. Recall that $D_i^{1,\text{out},S}$ denotes the face value of new secured debt, L_i its date-0 proceeds, and $\rho_i^{1,\text{out},S}(s)$ its repayment fraction in state s. The columns report:

$$\text{Debt price} \quad P_i^D = \frac{L_i}{D_i^{1,\text{out},S}} = \frac{1}{R}\sum_s \tilde{\pi}_s \rho_i^{1,\text{out},S}(s),$$

$$\text{Promised spread} \quad \left(\frac{1}{P_i^D} - R\right),$$

$$\text{Expected spread} \quad \frac{\sum_s \pi_s \rho_i^{1,\text{out},S}(s)}{P_i^D} - R = \frac{\sum_s \pi_s \rho_i^{1,\text{out},S}(s)}{P_i^D} - \frac{\sum_s \tilde{\pi}_s \rho_i^{1,\text{out},S}(s)}{P_i^D}.$$

The promised spread equals the expected loss per dollar invested plus the expected spread. The latter is the expected compensation for bearing credit risk, better known as the credit risk premium. Repayments are lower in bad states, which have a greater risk-neutral probability. Consequently, the price is low enough to offer a positive physical expected excess return.

For example, for Microsoft, lenders are promised a 42.28% return over the model period, or 37.28 percentage points above the risk-free return. Of this spread, 26.63 percentage points compensate for expected losses and 10.65 percentage points constitute the credit risk premium.

Table 5: New secured-loan prices and credit spreads.

Issuer	Debt price	Promised spread	Expected spread	Equity ($T)
Microsoft	0.703	37.28	10.65	3.04
Amazon	0.691	39.79	11.37	2.83
Google	0.687	40.56	11.59	4.61
Meta	0.672	43.89	12.54	1.70
Oracle	0.727	32.56	9.30	0.69
SpaceX / xAI / X	0.657	47.27	13.51	2.50
CoreWeave	0.952	0.00	0.00	0.06
Nebius	0.952	0.00	0.00	0.08
Landlord	0.838	14.30	4.09	0.85

Notes. The debt price is the amount advanced per dollar of promised repayment. The promised spread assumes full repayment; the expected spread accounts for losses using physical probabilities π_s. Both spreads are expressed in percentage points over the risk-free model-period return and are not annualized. Equity is post-issuance market value E_i^{MV}, including newly issued equity, expressed in trillions of dollars.

The credit risk premia in the table refer to new debt financing the buildout, rather than firms' legacy debt. This new debt is senior secured and non-recourse: lenders have priority over pledged cash flows from the newly financed assets but cannot claim the borrower's legacy assets or income. This distinction explains why Microsoft's new loans command a substantial risk premium while CoreWeave's command none. Microsoft's pledged project cash flows fall sharply in the low-productivity state, leaving its new lenders with a recovery of only 25.1%, even though its diversified legacy income allows it to honor ordinary corporate obligations in full. CoreWeave, by contrast, receives contractual payments from hyperscaler customers that remain solvent, including Microsoft, Meta and Alphabet. These pledged receipts suffice to repay its new secured lenders fully in every modeled state, even though its ordinary creditors recover only 43.0% in the low state. The difference therefore reflects the protection afforded to specific buildout loans by their collateral and contractual priority, rather than a ranking of Microsoft's and CoreWeave's overall corporate creditworthiness.

6 Financial Exposure and Downside Losses

The network model provides a framework for stress testing the AI buildout and tracing how losses propagate across firms and their investors. A decline in AI productivity reduces operating income, but its financial consequences depend on the obligations linking firms: one company's payment shortfall weakens its counterparties, potentially triggering further shortfalls and reducing the value of cross-held equity. Accounting for these spillovers allows us to assess both the scale of aggregate losses and their distribution across shareholders, lenders, and other creditors. We first stress the existing June 2026 network without additional buildout, then repeat the exercise with the equilibrium buildout to assess how the expansion changes financial exposure and downside losses. Finally, holding investment and promised contracts fixed, we isolate the contribution of contractual payment spillovers within the modeled network.

To compare promised exposures, define

$$\mathcal{D} = \frac{1}{R}\sum_{i\in\mathcal{I}}(D_i^S + D_i^J), \qquad \mathcal{W} = \frac{1}{R}\sum_{i,j\in\mathcal{I}} W_{ij}^J, \qquad \mathcal{M} = \sum_{i\in\mathcal{I}} E_i^{MV}. \tag{38}$$

Here $\mathcal{D}$ is total promised obligations, $\mathcal{W}$ is internal ordinary claims, and $\mathcal{M}$ is gross issuer equity market value.

To compare equity, debt and contractual claims consistently, define

$$Loss_c = P_c - \frac{X_c(\text{low})}{R}, \tag{39}$$

where P_c is the initial market value of claim portfolio c and $X_c(\text{low})$ its low-state payoff. This measures the valuation decline rather than the shortfall from promised face. In particular, risky lenders initially invest less than discounted face.

6.1 Stress Test Without the Buildout

Set $q_i^1 = 0$ for every compute user. All inherited cash, obligations and equity stakes remain in place, including uncommenced commitments signed by June 30, 2026. No new service or rental contracts, project loans or equity are issued. We clear and revalue the legacy network under the same three productivity states, holding fitted outside incomes $H_i^{0,MV}$ fixed at their benchmark levels.

Without additional buildout, promised obligations are $\mathcal{D} = \$5.22$ trillion, including $\mathcal{W} = \$2.00$ trillion of internal claims, against $\mathcal{M} = \$17.47$ trillion of revalued issuer equity. Obli-

gations equal $\mathcal{D}/\mathcal{M} = 29.8\%$ of equity and internal claims $\mathcal{W}/\mathcal{M} = 11.4\%$. The lower equity value reflects the removal of the investment opportunities associated with the buildout as well as the absence of new equity issuance.

The exposure measures above include claims held within the network, whereas Table 6 consolidates these internal holdings: eliminating cross-held equity reduces $\mathcal{M} = \$17.47$ trillion to \$17.147 trillion of outside-held equity. Similarly, $\mathcal{D} - \mathcal{W} \simeq \$3.22$ trillion measures outside obligations at promised face value discounted by R, whereas the table values these claims at their risk-adjusted market prices, giving \$1.168 trillion of lender debt and \$1.941 trillion owed to other outside creditors.

Table 6: Stress Test: No Buildout

Outside claimholder	Initial MV (\$T)	Low payoff/R (\$T)	Loss (\$T)
Outside-held equity	17.147	8.539	8.607
Outside lender debt	1.168	1.115	0.053
Other outside creditors	1.941	1.821	0.120
Consolidated total	20.255	11.475	8.780

The low state produces \$8.78 trillion of consolidated losses: \$8.61 trillion on outside-held equity, \$53 billion on lender debt and \$120 billion on other outside creditors.

In the bad state, OpenAI, Anthropic, Oracle and CoreWeave pay 20.94%, 32.97%, 31.45% and 65.35% of their ordinary obligations, respectively. The other firms pay ordinary creditors in full, and all inherited secured debt is repaid. Thus the legacy network is already exposed to AI risk, but most of the loss falls on shareholders rather than outside lenders.

6.2 Stress Test With the Buildout

We now go back to the benchmark model that features the additional 277.5 GW buildout, including its new contracts and financing.

With the equilibrium buildout, promised obligations rise to $\mathcal{D} = \$16.25$ trillion, comprising \$6.01 trillion of outside lender debt, $\mathcal{W} = \$8.24$ trillion of internal contracts, and \$2.01 trillion of other outside obligations. New debt accounts for \$4.80 trillion of lender promises, while new service and rental contracts account for \$3.68 trillion and \$2.56 trillion of internal promises, respectively. Against $\mathcal{M} = \$22.81$ trillion of post-issuance issuer equity, obligations equal $\mathcal{D}/\mathcal{M} = 71.2\%$ and internal claims $\mathcal{W}/\mathcal{M} = 36.1\%$, compared with 29.8% and 11.4%,

respectively, without additional buildout. Table 7 reports the corresponding market values of claims held by outside investors.

Table 7 reports low-productivity-state losses across claimholders. After consolidating internal holdings, the (discounted) payoff to outside investors falls from an initial market value of \$28.95 trillion to \$12.08 trillion, a loss of \$16.87 trillion, or 58.3%. Outside shareholders bear \$14.64 trillion of losses, outside lenders \$2.07 trillion, and other outside creditors \$0.16 trillion. Equity absorbs the largest dollar loss, but outside lenders still lose 42.3% of their initial investment.

The table also shows how internal holdings connect these losses across firms. Gross issuer equity, including internally held shares, falls from \$22.81 trillion to \$7.53 trillion in the low state. Across all equity and contractual claims, gross losses total \$21.25 trillion. This includes \$4.38 trillion of losses on internal assets that are already reflected in the payouts of the firms holding them. Eliminating these internal holdings removes the double counting and yields the consolidated loss of \$16.87 trillion.

The distribution of losses reflects both the decline in operating income and the transmission of payment shortfalls through the network. OpenAI, Anthropic and Oracle pay only 9.24%, 12.88% and 3.11% of their ordinary obligations, respectively, reducing suppliers' resources and the pledged receipts available to project lenders. CoreWeave pays 42.99% of ordinary claims while repaying its new secured loans in full, illustrating the protection provided by project collateral. Conversely, Microsoft and other solvent hyperscalers repay ordinary creditors in full while their non-recourse project lenders suffer losses: corporate solvency does not guarantee repayment from pledged project cash flows.

Shortfalls relative to promised repayments are larger than losses measured from initial market values, because risky claims are initially priced below discounted face value. Outside lenders' discounted repayment shortfall is \$3.18 trillion, of which \$3.05 trillion is on new loans, compared with their \$2.07 trillion market-value loss. Similarly, internal contractual promises go unpaid by \$5.75 trillion in discounted terms, whereas their loss from initial market value is \$3.74 trillion.

Relative to no additional buildout, consolidated losses rise by \$8.09 trillion, comprising \$6.04 trillion on outside equity, \$2.02 trillion on outside lender debt, and \$0.04 trillion on other outside claims. The aggregate increase equals $(\tilde{\theta} - \theta_{\text{low}})[a_V Q^1 - (b_V/2)(Q^1)^2]$, the valuation loss on the additional compute's operating benefits. The buildout therefore places substantially more capital at risk, while its financing and contractual links determine who

bears that risk.

Table 7: Stress Test: Benchmark Model

Claimholder / claim	Initial MV ($T)	Low payoff/R ($T)	Loss ($T)	% loss
Outside-held equity	22.135	7.493	14.642	-66.1%
Outside lender debt	4.893	2.824	2.069	-42.3%
Other outside creditors	1.922	1.760	0.163	-8.5%
Internally held equity	0.678	0.038	0.640	-94.4%
Internal debt/service/rent claims	6.226	2.489	3.736	-60.0%
Gross sum, before consolidation	35.855	14.605	21.250	-59.3%
Less internal holdings	6.904	2.528	4.376	-63.4%
Consolidated total	28.951	12.077	16.874	-58.3%

Notes. Loss equals initial market value minus the realized low-state payoff divided by the risk-free model-period return. Internal holdings are shown explicitly and eliminated only when computing the consolidated total; their losses are already transmitted into final outside claimant payouts. Internal holders are modeled firms. Equity includes incumbent and newly issued shares.

6.3 Isolating the Network Spillovers

To isolate the contribution of contractual payment spillovers, we ask how losses would change if firms received all payments promised by their counterparties within the modeled network, despite the low-productivity shock. We introduce an accounting backstop under which the government makes up shortfalls on internal contractual payments, preventing a customer's failure to pay from reducing its suppliers' receipts and triggering further payment shortfalls. Operating income remains at its low-state level, and investment, financing, promised contracts and ownership are held fixed at the buildout equilibrium. Equity holdings remain unprotected. Comparing this counterfactual with the actual network therefore separates the losses transmitted through missed contractual payments from those that remain when these payments are honored.

If ρ_j^G is issuer j's ordinary recovery from its own resources under the backstop, holder i of internal claim W_{ij}^J receives

$$W_{ij}^J \rho_j^G + G_{ij} = W_{ij}^J, \qquad G_{ij} = W_{ij}^J (1 - \rho_j^G). \tag{40}$$

Full receipts enter the holder's resources and, where pledged, its lenders' collateral. We clear the network again to determine actual obligor payments, equity and backstop transfers jointly. The backstop covers both legacy and new internal contracts. The exercise is an

accounting comparison rather than an ex-ante policy or a new investment equilibrium under a bailout.

Table 8: Stress test: No Internal Network Spillovers

Loss bearer	Actual network ($T)	Full receipts ($T)	Loss reduction ($T)
Outside shareholders	14.642	11.819	2.823
Outside lenders	2.069	-0.154	2.222
Other outside creditors	0.163	0.144	0.019
Private total	16.874	11.810	5.064
Accounting backstop	0.000	5.064	-5.064
Private plus backstop	16.874	16.874	0.000

The comparison in Table 8 shows that internal contractual payment spillovers are quantitatively substantial. Holding the buildout and low-state operating income fixed, suppressing these spillovers reduces private losses by $5.06 trillion, approximately 30% of the benchmark total. Outside shareholder losses fall by $2.82 trillion, outside lender losses by $2.22 trillion, and other outside creditor losses by $0.02 trillion.[7]

Thus, a sizable part of investors' downside exposure arises from the propagation of payment shortfalls through the network: customers' failure to pay weakens suppliers, reduces the receipts securing project loans, and transmits losses to those firms' creditors and shareholders.

7 Conclusion

By mid-2026, the AI network had already accumulated $5.22 trillion of obligations spanning financial debt, data-center leases, compute-service contracts, and IT hardware purchases. These exposures are substantial, but so is the equity supporting them: the modeled firms have $20.55 trillion of reference equity value. Large equity cushions and earnings from diversified legacy businesses protect many creditors and contractual counterparties, leaving shareholders to absorb most losses in the stress test without additional buildout. Yet aggre-

[7]Restoring service and rental receipts strengthens suppliers' resources and, where these receipts are pledged as collateral, improves project-debt recovery. Outside lenders consequently receive payoffs exceeding their initial market investment by approximately $0.15 trillion: their loans were initially priced to reflect repayment risk in the actual network, while the backstop raises their repayments in the low-productivity state.

gate financial strength can conceal weak links. Obligations are concentrated at particular firms, and equity elsewhere in the network is not automatically available to honor their promises. When these firms fail to pay, their shortfalls weaken suppliers, landlords, and other investors through contractual claims and cross-held equity.

The model provides a strong economic rationale for a large expansion of compute capacity. The optimal buildout features 179 GW of additional capacity, reflecting substantial expected productivity gains from AI. The decentralized equilibrium goes considerably further: 277.5 GW at a capital expenditure of $6.3 trillion, 55% above the social optimum. Firms capture the benefits of their own expansion without fully internalizing the reduction in other users' revenues as aggregate capacity grows. Limited liability can further encourage investment that shifts downside risk onto inherited creditors. Competitive, risk-based pricing of new contracts does not ensure an efficient scale of investment.

Financing this expansion makes the system considerably more vulnerable. Debt funds nearly 60% of new investment, raising aggregate lender debt relative to equity from 5.9% to 26.3%, with uneven increases across firms. New service and rental commitments further deepen counterparty exposures. Moreover, when the new secured debt is non-recourse, lenders depend on pledged project cash flows and cannot draw on the borrower's legacy businesses when those cash flows disappoint. In the low-AI-productivity state, consolidated losses reach $16.9 trillion, compared with $8.8 trillion without additional buildout. Shareholders lose $14.6 trillion, while outside lenders lose $2.1 trillion, or about 42% of their initial investment. The downside risk extends well beyond equity.

Internal network propagation is quantitatively important. Suppressing contractual payment spillovers reduces private losses by $5.06 trillion, approximately 30% of the total. The relationships that support the expansion also transmit its disappointments: missed customer payments erode suppliers' resources and the collateral backing project debt. The financial resilience of the AI buildout therefore depends on the scale of investment, the leverage used to finance it, and the location of repayment resources throughout the network.

References

Daron Acemoglu, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012. doi: 10.3982/ECTA9623.

Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015. doi: 10.1257/aer.20130456.

Levent Altinoglu. The origins of aggregate fluctuations in a credit network economy. *Journal of Monetary Economics*, 117:316–334, 2021. doi: 10.1016/j.jmoneco.2020.01.007.

Advait Arun. Bubble or nothing: Data center project finance. Report, Center for Public Enterprise, November 2025.

Eirik Eylands Brandsaas, Daniel Garcia, Robert Kurtzman, Joseph Nichols, and Adelia Zytek. Estimating aggregate data center investment with project-level data. Finance and Economics Discussion Series 2025-109, Board of Governors of the Federal Reserve System, 2025. URL https://www.federalreserve.gov/econres/feds/estimating-aggregate-data-center-investment-with-project-level-data.htm.

Anna M. Costello. Credit market disruptions and liquidity spillover effects in the supply chain. *Journal of Political Economy*, 128(9):3434–3468, 2020. doi: 10.1086/708736.

Larry Eisenberg and Thomas H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001. doi: 10.1287/mnsc.47.2.236.9835.

Andrea L. Eisfeldt and Adriano A. Rampini. Leasing, ability to repossess, and debt capacity. *The Review of Financial Studies*, 22(4):1621–1657, 2009. doi: 10.1093/rfs/hhn026.

Andrea L. Eisfeldt and Gregor Schubert. Generative AI and finance. *Annual Review of Financial Economics*, 17:363–393, 2025. doi: 10.1146/annurev-financial-112923-020503.

Andrea L. Eisfeldt, Gregor Schubert, Miao Ben Zhang, and Bledi Taska. Generative AI and firm values. *The Journal of Finance*, 2026. doi: 10.1111/jofi.70089. Advance online publication.

Matthew Elliott, Benjamin Golub, and Matthew O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–3153, 2014. doi: 10.1257/aer.104.10.3115.

Adam Feher, Emilia Garcia-Appendini, and Roxana Mihet. Is AI trained on public money? evidence from U.S. data centers. Working paper, January 2026.

Tom Fischer. No-arbitrage pricing under systemic risk: Accounting for cross-ownership. *Mathematical Finance*, 24(1):97–124, 2014. doi: 10.1111/j.1467-9965.2012.00526.x.

Jon Frost, Rudraksh Kansal, Kumar Rishabh, Vatsala Shreeti, and Leanne Si Ying Zhang. Circular relationships among AI firms. BIS Bulletin 137, Bank for International Settlements, October 2026. URL https://www.bis.org/publ/bisbull137.htm.

Antonio Gargano and Marco Giacoletti. Subsidizing the cloud: U.S. state incentives to data centers. Working paper, February 2026.

Pedro Gete and Amparo Mercader. Datacenter mortgages and originate-to-distribute. Working paper, February 2026.

Robin Greenwood and Samuel G. Hanson. Waves in ship prices and investment. *Quarterly Journal of Economics*, 130(1):55–109, 2015. doi: 10.1093/qje/qju035.

Michael C. Jensen and William H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976. doi: 10.1016/0304-405X(76)90026-X.

N. Gregory Mankiw and Michael D. Whinston. Free entry and social inefficiency. *RAND Journal of Economics*, 17(1):48–58, 1986. doi: 10.2307/2555627.

Ľuboš Pástor and Pietro Veronesi. Technological revolutions and stock prices. *American Economic Review*, 99(4):1451–1483, 2009. doi: 10.1257/aer.99.4.1451.

Adriano A. Rampini and S. Viswanathan. Collateral and capital structure. *Journal of Financial Economics*, 109(2):466–492, 2013. doi: 10.1016/j.jfineco.2013.03.002.

Phurichai Rungcharoenkitkul. The AI investment race. BIS Working Papers 1367, Bank for International Settlements, July 2026. URL https://www.bis.org/publ/work1367.htm.

René M. Stulz and Herb Johnson. An analysis of secured debt. *Journal of Financial Economics*, 14(4):501–521, 1985. doi: 10.1016/0304-405X(85)90024-8.

Stijn Van Nieuwerburgh. Financing the AI buildout. Brookings Papers on Economic Activity, Fall 2026 conference draft, 2026. URL https://www.brookings.edu/wp-content/uploads/2026/09/4c_Van-Nieuwerburgh.pdf.

A Model Inputs: Data, Sources and Construction

This appendix lists every input of the network model at its reference date, June 30, 2026: the value the model reads, where it comes from and the rule that turns the source figure into that value.

A.1 Sources

Table A1: Sources and balance-sheet dates.

Company	Balance sheet	Filings	Other sources
NVIDIA	Apr 26, 2026	10-Q Q2 FY2026; 10-Q Q3 FY2026; Q4 FY2026 earnings release; 10-Q Q1 FY2027; 10-Q Q2 FY2027	Epoch AI, Data on AI chip owners; CoreWeave 10-Q Q1 2026; CoreWeave 10-Q Q2 2026; Nebius 6-K Q1 2026; Nebius 6-K Q2 2026
Microsoft	Jun 30, 2026	10-K FY2026	Microsoft Fiscal Year 2026 Second Quarter Earnings Conference Call (Jan. 28, 2026) Reuters (Sept. 29, 2026), Anthropic cloud commitments
Amazon	Jun 30, 2026	10-Q Q2 2026; 8-K, February 2026	Reuters (Sept. 29, 2026), Anthropic cloud commitments
Alphabet	Jun 30, 2026	10-Q Q2 2026	Reuters via BNN Bloomberg (Aug. 22, 2025), Meta signs over US$10 billion cloud deal with Google, source says Reuters (Sept. 29, 2026), Anthropic cloud commitments

continued on the next page

Company	Balance sheet	Filings	Other sources
Meta	Jun 30, 2026	10-Q Q2 2026	CoreWeave 8-K, September 25, 2025; CoreWeave 8-K, March 31, 2026; Nebius 20-F FY2025; Nebius 6-K, March 2026 Bloomberg (Sept. 19, 2025), Oracle in Talks With Meta on $20 Billion AI Cloud Deal Oracle Q2 2026 Earnings Call Transcript (Dec. 10, 2025, via Insider Monkey)
Oracle	May 31, 2026	10-K FY2026	The Wall Street Journal (Sept. 10, 2025), Oracle, OpenAI Sign $300 Billion Cloud Deal Data Center Dynamics (Sept. 11, 2025), OpenAI signs $300bn cloud deal with Oracle – report
CoreWeave	Jun 30, 2026	10-Q Q2 2026; 10-Q Q1 2025; 8-K, September 25, 2025; 8-K, March 31, 2026	NVIDIA Schedule 13G/A, January 2026 CoreWeave Expands Agreement with OpenAI by up to $6.5B (Sept. 25, 2025) Tech Insider (May 10, 2026), CoreWeave's Anthropic Deal: 12% Surge, 6.8B Backlog [2026]
Nebius	Jun 30, 2026	6-K Q2 2026; 20-F FY2025; 6-K, March 2026	NVIDIA Schedule 13G

continued on the next page

Company	Balance sheet	Filings	Other sources
OpenAI	Jun 30, 2026	Microsoft 10-K FY2026; Amazon 10-Q Q2 2026; Amazon 8-K, February 2026; CoreWeave 10-Q Q1 2025	The Information (June 16, 2026), OpenAI Burned $3.7 Billion in First Three Months of 2026 The Information (June 22, 2026), How OpenAI's Web of Business Relationships Could Complicate Its IPO OpenAI (Mar. 31, 2026), OpenAI raises $122 billion to accelerate the next phase of AI Bloomberg Law (Mar. 31, 2026), OpenAI Valued at $852 Billion After Mega $122 Billion Round CoreWeave (Sept. 25, 2025), expanded OpenAI agreement
Anthropic	Jun 30, 2026	Form S-1, October 26, 2026; Amazon 10-Q Q2 2026; Alphabet 10-Q Q2 2026; Microsoft 8-K, July 29, 2026	Reuters via Investing.com (Sept. 29, 2026), Anthropic's $518 billion AI buildout hinges largely on deals that cannot be canceled, filing shows
SpaceX/xAI/X	Jun 30, 2026	10-Q Q2 2026; S-1/A, June 2026	S&P Global Compustat (WRDS); NVIDIA 10-Q Q2 FY2027

A.2 Legacy Balance Sheets

Table A2 covers eleven firms and includes obligations to creditors inside and outside the network. Appendix A.5 discusses the landlord balance sheet.

Panel A reports cash, secured obligations, and ordinary obligations. Panel B breaks out the ordinary obligations into commenced and uncommenced obligations.

We assume OpenAI had $100 billion in cash on June 30, 2026. OpenAI held over $73 billion in cash and marketable securities at March 31, 2026 (*The Next Web*, June 17, 2026, citing *The Information*).

We assume that all of OpenAI's purchase commitments are to Microsoft, Amazon, Alphabet, Oracle and CoreWeave. Together, these contracts total $761.4 billion in nominal terms, or $659.3 billion in present value.

Microsoft reports $194.060 billion in purchase commitments and $34.566 billion in construction commitments as of June 30, 2026. The latter mainly cover data center buildings and improvements. Applying 0.8659 to the combined nominal amount gives the $198.0 billion PV shown in the table.

Anthropic's Form S-1 reports $321.4 billion in uncommenced leases ($222.4 billion in PV), all assigned to outside lessors. The $265.6 billion in purchase commitments has a PV of $230.0 billion, of which $224.5 billion is owed to named providers inside the network and $5.5 billion to outside suppliers.

CoreWeave's secured debt includes delayed-draw term loans (including non-recourse DDTL 4.0) and OEM equipment financing, net of issuance costs. Its unsecured debt includes senior and convertible notes, software-license financing and the Magnetar loan.

SpaceX/xAI/X's secured financings are backed by machinery and equipment and include failed sale-leaseback obligations.

Table A2: Cash and legacy obligations on June 30, 2026 (PV, $ billions).

Panel A. Cash, secured obligations and ordinary obligations

Firm	Cash C_i^0	Secured D_i^{S0}	Ordinary D_i^{J0}
NVIDIA	13.2	0.0	182.5
Microsoft	20.9	0.0	596.9
Amazon	78.2	0.0	597.0
Alphabet	55.9	0.0	906.3
Meta	15.5	0.0	633.2
Oracle	31.3	0.0	386.3
CoreWeave	5.5	18.2	61.6
Nebius	8.0	0.0	24.2
OpenAI	100.0	0.0	666.0
Anthropic	112.4	0.0	454.7
SpaceX/xAI/X	93.5	13.4	41.9

Panel B. Composition of ordinary obligations

Firm	Unsecured debt	Commenced leases	Payables	Uncommenced leases	Purchase and Construction Commitment	Total D_i^{J0}
NVIDIA	8.5	4.3	13.1	22.4	134.2	182.5
Microsoft	40.3	88.5	42.4	227.7	198.0	596.9
Amazon	132.2	109.8	147.4	95.0	112.6	597.0
Alphabet	100.2	20.6	20.3	63.0	702.2	906.3
Meta	83.7	38.1	15.9	193.1	302.5	633.2
Oracle	129.5	37.9	11.0	179.9	28.0	386.3
CoreWeave	16.9	16.5	3.6	24.6	0.0	61.6
Nebius	8.5	1.6	1.1	8.3	4.6	24.2
OpenAI	0.0	0.8	6.0	—	659.3	666.0
Anthropic	0.0	1.5	0.8	222.4	230.0	454.7
SpaceX/xAI/X	25.0	1.4	8.2	0.0	7.2	41.9

Notes. We sum each firm's nominal commitments across maturities by category. To compute the present values, we multiply the uncommenced lease totals by 0.692 (15-year annuity at 5% discount rate) and the purchase and construction commitment totals by 0.8659 (5-year annuity at 5% discount rate). The data sources for each firm are listed in Table A1.

A.3 Claims Between Firms in the Named Network

The 42 claims between modeled firms are ordinary legacy obligations $W_{ij}^{0,J}$.

Section A.6 gives the calculation for all eight NVIDIA receivables.

We assume the modeled landlord is owed 90% of commenced lease PV and 95% of uncommenced lease PV for Microsoft, Alphabet, Meta, Oracle, CoreWeave and Nebius. The corresponding shares are 35% and 75% for Amazon, 80% and 95% for NVIDIA, and 0% for both categories for OpenAI and Anthropic. For SpaceX/xAI/X we assign 80% of commenced lease PV to the landlord. No uncommenced leases are recorded for SpaceX/xAI/X. The balance is owed to outside lessors. Section A.5 uses these claims to calculate rented capacity.

Table A3: Named claims between modeled firms.

Holder	Obligor	Type	U/C*	Nominal	PV	Source
Microsoft	OpenAI	purchase	U	281.0	243.3	45% of commercial backlog (Microsoft earnings call, Jan. 2026)
Microsoft	OpenAI	account payable	C	6.0	6.0	Receivable from OpenAI (Microsoft FY2026 10-K, p. 60)
Amazon	OpenAI	purchase	U	138.0	119.5	AWS contract and expansion (Amazon Q2 10-Q)
Alphabet	OpenAI	purchase	U	20.0	17.3	Google Cloud capacity, amount assumed (Alphabet earnings call)
Oracle	OpenAI	purchase	U	300.0	259.8	Five-year cloud contract starting in 2027 (WSJ, Sept. 2025)
CoreWeave	OpenAI	purchase	U	22.4	19.4	Three compute orders, combined ceiling (CoreWeave, Sept. 2025)
Microsoft	Anthropic	purchase	U	31.4	27.2	Minimum Azure purchases (Reuters, Sept. 2026)
Amazon	Anthropic	purchase	U	110.0	95.2	Minimum AWS purchases (Reuters, Sept. 2026)
Alphabet	Anthropic	purchase	U	111.1	96.2	Minimum Google Cloud purchases (Reuters, Sept. 2026)

* C denotes commenced obligations (accounts payable and leases already in effect). U denotes uncommenced obligations (signed purchase commitments still to be fulfilled and leases yet to begin).

continued on the next page

Holder	Obligor	Type	U/C*	Nominal	PV	Source
CoreWeave	Anthropic	purchase	U	6.8	5.9	MarketMinute report (contract value estimate cited by Tech Insider)
CoreWeave	Microsoft	purchase	U	10.0	8.7	Compute capacity through 2030 (The Information, Nov. 2024)
Nebius	Microsoft	purchase	U	17.4	15.1	Capacity fees payable regardless of use, subject to delivery (Nebius 20-F)
CoreWeave	Alphabet	purchase	U	5.0	4.3	Capacity for Google Cloud, amount assumed (Reuters, June 2025)
Alphabet	Meta	purchase	U	10.0	8.7	Six-year Google Cloud contract (Reuters, Aug. 2025)
Oracle	Meta	purchase	U	20.0	17.3	Cloud contracts confirmed, amount assumed (Oracle call and Bloomberg)
CoreWeave	Meta	purchase	U	35.2	30.5	Two capacity orders, combined ceiling (CoreWeave 8-Ks)
Nebius	Meta	purchase	U	14.9	12.9	Initial order and capacity expansion (Nebius 20-F and 6-K)
NVIDIA	Microsoft	account payable	C	7.5	7.5	Chip receivable allocated by GPU share (NVIDIA 10-Q and Epoch AI)
NVIDIA	Amazon	account payable	C	3.3	3.3	Chip receivable allocated by GPU share (NVIDIA 10-Q and Epoch AI)
NVIDIA	Alphabet	account payable	C	2.8	2.8	Chip receivable allocated by GPU share (NVIDIA 10-Q and Epoch AI)
NVIDIA	Meta	account payable	C	4.8	4.8	Chip receivable allocated by GPU share (NVIDIA 10-Q and Epoch AI)
NVIDIA	Oracle	account payable	C	2.8	2.8	Chip receivable allocated by GPU share (NVIDIA 10-Q and Epoch AI)
NVIDIA	CoreWeave	account payable	C	3.5	3.5	Chip receivable estimated from equipment additions (NVIDIA and CoreWeave 10-Qs)
NVIDIA	Nebius	account payable	C	0.7	0.7	Chip receivable estimated from server and network additions (NVIDIA 10-Q and Nebius 6-Ks)

* C denotes commenced obligations (accounts payable and leases already in effect). U denotes uncommenced obligations (signed purchase commitments still to be fulfilled and leases yet to begin).

continued on the next page

Holder	Obligor	Type	U/C*	Nominal	PV	Source
NVIDIA	SpaceX/xAI/X	account payable	C	2.9	2.9	Chip receivable estimated from reclassified revenue (NVIDIA Q2 10-Q)
Landlord	NVIDIA	leases	C	3.5	3.5	Operating leases (Q1 10-Q)
Landlord	NVIDIA	leases	U	30.8	21.3	Future leases, mainly for R&D data centers (Q1 10-Q)
Landlord	Microsoft	leases	C	79.7	79.7	Operating and finance leases (FY2026 10-K, Note 13)
Landlord	Microsoft	leases	U	312.6	216.4	Future lease payments (FY2026 10-K, Note 13)
Landlord	Amazon	leases	C	38.4	38.4	Operating and finance leases (Q2 10-Q)
Landlord	Amazon	leases	U	102.9	71.2	Future lease payments (Q2 10-Q)
Landlord	Alphabet	leases	C	18.6	18.6	Operating and finance leases (Q2 10-Q)
Landlord	Alphabet	leases	U	86.5	59.8	Future leases, including short-term leases (Q2 10-Q)
Landlord	Meta	leases	C	34.3	34.3	Operating leases and a finance lease asset proxy (Q2 10-Q)
Landlord	Meta	leases	U	265.0	183.4	Future leases, including Hyperion (Q2 10-Q)
Landlord	Oracle	leases	C	34.1	34.1	Operating and finance leases (FY2026 10-K)
Landlord	Oracle	leases	U	247.0	170.9	Future leases, mainly data centers (FY2026 10-K)
Landlord	CoreWeave	leases	C	14.9	14.9	Operating and finance leases (Q2 10-Q)
Landlord	CoreWeave	leases	U	33.7	23.3	Future lease payments (Q2 10-Q)
Landlord	Nebius	leases	C	1.5	1.5	Operating leases (Q2 6-K)
Landlord	Nebius	leases	U	11.5	7.9	Future leases on data centers and equipment (Q2 6-K)
Landlord	SpaceX/xAI/X	leases	C	1.1	1.1	Current operating and finance lease obligations (Q2 10-Q)

* C denotes commenced obligations (accounts payable and leases already in effect). U denotes uncommenced obligations (signed purchase commitments still to be fulfilled and leases yet to begin).

A.4 Equity Market Values and Equity Stakes

The observed equity market value on June 30, 2026, E_i^{obs}, is an input in our procedure to back out legacy outside assets $H_i^{0,MV}$. Table A4 reports E_i^{obs} in billions of dollars, and explains how we compute it. For publicly-listed companies this is the closing stock price times the number of shares outstanding of all share classes (incl. unlisted share classes).

For OpenAI and Anthropic, who are still private companies at the time of this writing, we use the post-money valuation of the latest fund raise prior to June 30, 2026. Those data are from Pitchbook.

We derive the landlord sector's equity market value from property values and secured debt, as explained in Appendix A.5.

Table A4: Reference market value of equity at June 30, 2026, billions of dollars.

Firm	$E_i^{0,MV}$	Basis	Source
NVIDIA	4,846.4	CRSP market capitalization	CRSP via WRDS [Market]
Microsoft	2,771	CRSP market capitalization	CRSP via WRDS [Market]
Amazon	2,564	CRSP market capitalization	CRSP via WRDS [Market]
Alphabet	4,348.3	12,230m Class A, B and C shares at the June 30 class prices (357.37, 353.33)	10-Q cover; CRSP, Yahoo Finance [Filing]
Meta	1,435.3	2,548m Class A and B shares × 563.29	10-Q; CRSP [Filing]
Oracle	422.1	CRSP market capitalization	CRSP via WRDS [Market]
CoreWeave	54.9	551m Class A and B shares × 99.54	10-Q; CRSP [Filing]
Nebius	80.9	292.9m shares, NVIDIA's pre-funded warrants included, × 276.17	6-K; CRSP [Filing]
OpenAI	852	Post-money valuation of the March 2026 round	OpenAI announcement, Mar 31, 2026; Bloomberg Law [Reported]
Anthropic	739.9	12,562m common-equivalent shares, including the shares under Amazon's notes, × 58.90 (Series H)	S-1 [Filing]
SpaceX/xAI/X	2,251.3	13,176m Class A and B shares × 170.86	10-Q; Compustat close [Filing]
Landlord sector	188.3	Property value 821.23 less debt 632.95	Derived [Derived]

Table A5 reports ownership shares of common-equivalent equity that modeled firms have in

each other's companies, O^0_{ij}.

Table A5: Equity stakes between modeled firms.

Holder	Issuer	O_{ij}	Source and calculation
Alphabet	Anthropic	11.88%	Alphabet 10-Q Q2 2026; Anthropic S-1. We assume the remeasured holding of 87.9 is Anthropic and value it at the Series H price: 87,900/58.90095/12,562.426 = 11.88%. [Derived; assumption]
Amazon	Anthropic	27.15%	Amazon 10-Q Q2 2026; Anthropic S-1. We assume the preferred is valued at the Series H price and the notes fully convert, including waiver of the ownership cap: (92,500/58.90095 + 1,840.445)/12,562.426 = 27.15%. [Derived; assumption]
Amazon	OpenAI	3.37%	Amazon 10-Q Q2 2026 Calculation: 28.7 / 852 = 3.37%. [Filing]
Microsoft	Anthropic	1.60%	We assume the retained estimated holding value is 11.842105: 11.842105/739.94 = 1.6004%, rounded to 1.60%. [Assumption]
Microsoft	OpenAI	25.00%	Microsoft 10-K FY2026, p. 60 Calculation: As-converted interest stated in the 10-K. [Filing]
NVIDIA	CoreWeave	8.57%	Schedule 13G/A, Jan 2026; CoreWeave 10-Q Calculation: 47.213 / 551 = 8.57%. [Filing]
NVIDIA	Nebius	7.60%	Schedule 13G; Nebius 6-K Calculation: (1.190 + 21.066) / (238.400 + 33.455 + 21.066) = 22.256 / 292.921 = 7.60%. [Filing]
NVIDIA	OpenAI	3.52%	Bloomberg Law, Mar 31, 2026. We assume the reported 30 investment is fully funded and held at June 30: 30/852 = 3.52%. [Reported; assumption]
OpenAI	CoreWeave	1.63%	CoreWeave 10-Qs. We assume all 9m shares issued to OpenAI remain held at June 30: 9/551 = 1.63%. [Filing; assumption]
Alphabet	SpaceX/xAI/X	4.18%	Alphabet 10-Q Q2 2026; SpaceX 10-Q Calculation: 94.1 / 170.86 = 551m shares; 551 / 13,176 = 4.18%. [Filing]

A.5 Data-center Capacity and the Landlord Sector

Table A6 covers the data-center capacity of the six hyperscalers. We use the September 12, 2026 Cleanview extract to sum the facility capacity of US projects that are operating or under construction, regardless of opening year. Projects whose developer is Microsoft, Amazon, Google, Meta, Oracle or xAI count toward that hyperscaler's owned capacity.

We assign projects by other developers to the landlord sector. We assume neoclouds lease all real estate.[8]

Table A6: Owned and rented data-center capacity.

Firm	Owned US	Owned	Owned $K_j^{0,RE}$	Rent PV	Rented	ϕ_j
	GW	GW	\$B	\$B	GW	%
Microsoft	8.7	10.5	136	296.0	18.3	36.40
Amazon	21.9	26.3	342.1	109.6	6.8	79.54
Alphabet	13.6	16.3	212.5	78.4	4.8	77.15
Meta	11	13.2	171.4	217.7	13.4	49.50
Oracle	0.0284	0.0341	0.4	205.0	12.7	0.27
SpaceX/xAI/X	1.9	2.2	29.1	1.1	0.1	96.95

Owned and rented capacity. The six hyperscalers account for 57.146 GW in the US. Multiplying by 1.2 gives their 68.575 GW worldwide capacity. At \$13 million per MW, its real estate and power infrastructure have a development cost of $K_i^{0,RE}$ billions of dollars, reported in the third column.

From rent to rented GW. We add the PVs of each tenant's commenced and uncommenced lease obligations to the modeled landlord (Table A3). The 15-year annuity factor at 5% is 10.3797. Annual gross rent per GW is 12% of the \$13 billion development cost, or \$1.56 billion. We use these quantities to calculate rented capacity and the owned-site fraction.

$$\text{rented GW}_i = \frac{\text{rent PV}_i}{10.3797 \times 1.56} = \frac{\text{rent PV}_i}{16.1923}, \qquad \phi_i = \frac{\text{Owned}_i}{\text{Owned}_i + \text{Rented}_i}.$$

Landlord balance sheet. The landlord sector aggregates third-party owners of data-center capacity. Table A7 gives its estimated balance sheet at the reference date.

[8]Cleanview lists 0.052 GW in data-center capacity for CoreWeave and 0.200 GW for Nebius, all under construction.

Table A7: Landlord balance sheet at June 30, 2026 ($ billions).

Assets		**Liabilities and equity**	
Cash	0.00	Secured debt	632.95
Data-center property	821.23	Equity	188.28
Total assets	821.23	Total liabilities and equity	821.23

Capacity and development cost. After excluding projects developed by the six hyperscalers, 924 US projects are operating or under construction, with 50.72 GW of reported capacity. Scaling US capacity by 1.20 to obtain global datacenter capacity for every firm and the landlord sector gives 60.86 GW worldwide. At a cost of $13 billion per GW (comprised of a $11 million per MW development cost for the shell and core plus $2 million per MW for the power infrastructure) the development cost is $B = \$791.19$ billion.

Debt and equity. We assume zero cash. Outside lenders hold the secured debt, which we set to $D^{S0} = 0.80B$ using the same loan-to-cost assumption as for the buildout. Annual gross rent is $0.12B$ and operating expenses are $0.02B$. We value the property by multiplying annual net rent of $0.10B$ by the 15-year annuity factor of 10.3797.[9] Reference equity is the property value less secured debt.

A.6 NVIDIA's receivables by end customer

NVIDIA generally sells accelerators through equipment manufacturers. We attribute the receivables to the end customers by multiplying their estimated purchases in the quarter ended April 26, 2026 by $K = 0.5610$. We compute K as the simple average of quarter-end net accounts receivable divided by that quarter's revenue over four NVIDIA fiscal quarters, from Q2 FY2026 through Q1 FY2027 (filings in Table A1).

Applying K to NVIDIA's originally reported Hyperscale revenue of $37.869 billion gives $21.245 billion in estimated receivables from five hyperscalers. We allocate this amount in proportion to their October–December 2025 NVIDIA chip additions in Epoch AI, measured in H100 equivalents. CoreWeave's $6.190 billion and Nebius's $1.278 billion are additions to technology equipment and to server and network equipment, respectively. We prorate their additions by the number of days each calendar quarter overlaps NVIDIA's fiscal quarter and assume these purchases were NVIDIA systems. For SpaceX/xAI/X, we use the $5.181 billion in April-quarter revenue that NVIDIA later reclassified as Hyperscale, assuming the customer was xAI.

[9] The 15-year annuity factor at 5% is $[1 - 1.05^{-15}]/0.05$.

Table A8: Estimated purchases and NVIDIA receivables by end customer ($ billions).

End customer	Estimated purchases	Receivable PV
Microsoft	13.45	7.55
Amazon	5.86	3.29
Alphabet	5.01	2.81
Meta	8.49	4.76
Oracle	5.05	2.83
CoreWeave	6.19	3.47
Nebius	1.28	0.72
SpaceX/xAI/X	5.18	2.91
Total	**50.52**	**28.34**

These eight claims account for $28.34 billion of NVIDIA's $40.71 billion in net receivables at April 26, 2026.

A.7 Lab procurement and subcontracting shares

Lab procurement shares (ω). We assume each lab allocates all its new compute purchases among modeled operators in proportion to the legacy compute claims each operator holds on that lab. Let v_{ja} be the present value of operator j's claims on lab a in Table A3. The shares in Table A9 are calculated as

$$\omega_{ja} = \frac{v_{ja}}{\sum_{j'} v_{j'a}}.$$

Table A9: Lab procurement shares and present values of legacy compute claims ($ billions).

Operator	Claims on OpenAI (PV)	$\omega_{j,OpenAI}$	Claims on Anthropic (PV)	$\omega_{j,Anthropic}$
Microsoft	249.3	37.47%	27.2	12.11%
Amazon	119.5	17.96%	95.2	42.42%
Alphabet	17.3	2.60%	96.2	42.85%
Oracle	259.8	39.05%	0.0	0.00%
CoreWeave	19.4	2.92%	5.9	2.62%
Total	**665.3**	**100%**	**224.5**	**100%**

Microsoft's claim on OpenAI includes OpenAI's purchase commitment and its separate $6.0 billion payable to Microsoft. Anthropic's $224.5 billion equals $0.8659 \times (31.4 + 110.0 + 111.1 + 6.8)$ billion for Microsoft, Amazon, Alphabet and CoreWeave, respectively.

Subcontracting shares (β). For hyperscaler h, let S_{kh} be its compute purchases from neocloud k (Table A3), L_h its compute sales to OpenAI and Anthropic, and I_h the imputed value of its own compute use. All three are present values.

$$\underbrace{\frac{S_{kh}}{L_h}}_{\text{lab sales only}} \quad (L_h > 0), \qquad \underbrace{\frac{S_{kh}}{L_h + I_h}}_{\text{lab sales and own use}} \quad = \beta_{kh}.$$

Table A10 reports both ratios. When $L_h > 0$, the second equals the first times the labs' share of demand, $L_h/(L_h + I_h)$. We solve for own use and β jointly.

Meta has no recorded sales to either lab, so its lab-sales ratio is undefined (n.a.). Its second ratio uses the value of all 40 GW of own use as the denominator. Alphabet's CoreWeave

contract amount is undisclosed. We use the mean of Microsoft's two lab-sales ratios (4.38%) and multiply by the labs' share of Alphabet's demand (53.6%).

Table A10: Neocloud subcontracting ratios with and without the buyer's own compute use.

Hyperscaler buyer	Neocloud supplier	Lab sales only (%)	Lab sales and own use (%)
Microsoft	CoreWeave	3.20	2.22
Microsoft	Nebius	5.57	3.87
Meta	CoreWeave	n.a.	3.37
Meta	Nebius	n.a.	1.42
Alphabet	CoreWeave	4.38	2.35

A.8 AI Shares of Legacy Outside Assets

Outside income in state s is proportional to $1 - \alpha_i + \alpha_i \theta_s$. The values for α_i are the rough shares of each firm's revenue accounted for by AI-related business activities. Table A11 justifies the choices for α_i, while Table A12 reports the cash flows on the legacy outside assets in each of the three states of the world in period 1, $X_i^{H,0}$.

Table A11: AI exposure of outside income.

Firm	α_i	Basis
NVIDIA	0.9	Data-center revenue is about 90% of NVIDIA's total.
Microsoft	0.35	AI-related cloud and software revenue as a share of the total.
Amazon	0.25	AWS and AI-related revenue as a share of a retail-dominated total.
Alphabet	0.3	AI-exposed share of a search- and advertising-dominated total.
Meta	0.25	AI-exposed share of an advertising-dominated total.
Oracle	0.5	OCI and AI-related revenue as a share of the total.
CoreWeave	1	A pure AI cloud.
Nebius	0.9	An AI cloud with small other businesses.
OpenAI	1	A pure AI lab.
Anthropic	1	A pure AI lab.
SpaceX/xAI/X	0.55	xAI and X within a consolidated group whose launch and Starlink businesses are not AI.
Landlord sector	0	Rents are contractual; the sector's outside income is state-independent.

Table A12: Legacy outside-income payoffs $X_i^{H,0}(s)$

Firm	α_i	High	Middle	Low
OpenAI	1.00	1615.26	828.34	41.42
Anthropic	1.00	1536.36	787.88	39.39
Microsoft	0.35	4144.57	3110.37	2076.17
Amazon	0.25	3576.21	2889.87	2203.53
Google	0.30	6842.71	5325.07	3807.42
Meta	0.25	2342.66	1893.06	1443.45
Oracle	0.50	54.51	36.96	19.40
CoreWeave	1.00	141.79	72.71	3.64
Nebius	0.90	162.30	87.49	12.69
NVIDIA	0.90	10741.26	5790.44	839.61
SpaceX / xAI / X	0.55	3463.81	2275.08	1086.35
Landlord	0.00	-86.00	-86.00	-86.00