

Yale Center Beijing and Ted and Logan Webinars on High-Tech Industries

Current Events (NVDA Earnings, AI Open vs Closed, Anthropic IPO) & YCB Program Preview

Edward A. Snyder and Logan Bender

2 September 2026

som.yale.edu/hightech

The YCB Program: Analysis and Valuation of High-Tech Companies

Yale Center Beijing · Yale School of Management · October 9 – November 20, 2026

8

Live sessions

90 minutes each

Oct 9 – Nov 20

2026

across Q3 earnings season

~30

Companies covered

China + Magnificent 7

Program fees

*Stay abreast of our social media channels for
info and updates.*

Yale SOM Certificate

Awarded upon completion — attendance and participation in all eight sessions.

Faculty

Logan Bender, CFA

Yale MBA · Global Equity Research Analyst

Co-developer of the most subscribed Global Network for Advanced Management elective (2020–2024). Expert in the analysis and valuation of China and U.S. technology companies.

Prof. Edward A. Snyder

William S. Beinecke Professor of Economics and Management

Former Dean, Yale SOM. Authority on industrial organization and competitive strategy.

The YCB Program: Agenda, Format and Coverage

Agenda

#	Date	Session topic
S1	Oct 9	Program Kickoff, Q3 Preview & ER Framework
S2	Oct 10	The Magnificent 7 & China Leaders
S3	Oct 16	First Insights — early Q3 results
S4	Oct 23	Valuation Primer
S5	Oct 30	IPO Timing & Catalysts: Anthropic & OpenAI
S6	Nov 6	Semis & Software: AMD, QCOM, PLTR & AI Infra
S7	Nov 13	China Tech: Tencent, Alibaba, Baidu
S8	Nov 20	Wrap-Up & Synthesis: Portfolio Positioning

Coverage Universe Includes

China: Alibaba, JD.com, Meituan, PDD, Tencent, Baidu
Magnificent 7: Alphabet, Amazon, Apple, Nvidia, Tesla, Microsoft, Meta

Format

- Eight live 90-minute sessions across a full quarterly earnings cycle
- ~5 hrs/week valuation lab — readings, cases, modeling
- Office hours after sessions 2, 3, 5 and 7

Source: Yale Center Beijing · Yale School of Management.

China-US Rivalry: Role of Open-Weight Models

- Core components are publicly released, allowing anyone to download it
- DeepSeek was the first major open-weight model to capture the attention
- Best Chinese models are open-weights (DeepSeek, Qwen, Kimi, GLM)
- Best U.S. models are proprietary (Claude, GPT, Gemini)
- U.S. hyperscalers developed some open-weight models (LLaMA, Gemma, Nemotron) but they don't perform as well
- Open-weight models are generally cheaper to operate and allow easy customization

Background Insight: Open versus Closed Systems


- For decades, high-tech industries have featured a mix of open and closed systems: Windows vs. Linux, STATA vs. R, iOS vs. Android
- Most proprietary systems have significant open-source components
- Many ‘open’ systems need extra proprietary components to work (Android is open-source, but necessary device drivers are not)
- Both models coexist and complement each other
- Why would companies spend millions of dollars developing valuable LLMs that others can modify?

NVIDIA News

- Earnings for Q2 released August 26
- Announcement of “pause” in revenue-sharing deals (August 27)
- Also in August, released Nemotron 3.5 Lightning, a “lightweight” version of its open weight model
- Planning to spend \$6 billion to build “one of the world’s most powerful open-weight AI models” to compete with DeepSeek, Qwen et al.
- Upcoming slides: unpacking the August 26 earnings event — what beat, what was forecast, and what stands behind it

NVIDIA: Another Beat and Raise

July quarter — reported vs. consensus, and what drove the gap



Reported beat consensus on both lines; Data Center is ~93% of revenue.

What happened

- **Total revenue \$96.2B**, +18% Q/Q, against a \$92.1B Street number — a beat of roughly \$4B.
- **Data Center \$89.0B**, also +18% Q/Q, versus \$85.0B expected. This is the whole story of the quarter.
- **Two drivers:** stronger inference demand — not just training — and the Blackwell ramp reaching scale.
- **Beat and raise.** Q3 guidance of \$108B against a \$104.2B Street number. The guide moved up with the print, which is what separates a good quarter from a re-rating.

+18% Q/Q


Sequential growth

Total revenue and Data Center alike

Source: Jefferies, “NVIDIA — Is US\$1T Within Reach?” Consensus = sell-side mean as compiled by Jefferies. \$ in billions.

Is US\$1 Trillion Within Reach?

The new FY28 outlook sits far above the Street — and some analysts call it conservative



FY28 guide implies +70% Y/Y — about \$130B above consensus.

The bridge to \$1T

- **The raise.** About \$700B of revenue in FY28, up 70% — roughly \$130B more than Wall Street expected.
- **A first.** NVIDIA has never forecast a full year before. It still gave the usual quarter — \$108B for Q3 — and added the year on top.
- **Why it can.** Each new generation earns more revenue from the same amount of electricity, so the same buildout is worth more.

\$18B → \$25B → \$40B

Revenue per gigawatt

Hopper → Blackwell → Vera Rubin

Source: NVIDIA Q2 FY27 earnings call, Aug. 26, 2026; sell-side research. FY27E implied by the FY28 outlook of ~\$700B at +70% Y/Y. FY29 is an analyst path, not company guidance. \$ in billions.

What Underwrites the Guide

The binding constraint

Supply, not demand

Customers want more than 70%; 70% is what NVIDIA can build

What is actually scarce

Memory

Not NVIDIA's own chips — the memory that goes around them

Committed to suppliers

\$279B

More than doubled — cash pledged to lock up parts

How to read it

- **The 70% is a supply number, not a demand number.** Management said demand is much greater than 70% — customer forecasts point to it doubling. Seventy percent is simply what they can build, and they expect to stay sold out into 2028.
- **That is why they can forecast a year out.** When you are sold out, next year's revenue is easier to see than it sounds.
- **The \$279B is the proof — and the risk.** It locks up the parts to build through next year. It is also a fixed bill if demand slows.

From the Call: What Management Actually Said

The demand story, the money question, and one number that matters in this room

Compute per agent

15 – 100x

Versus a human doing the same work

Neocloud capacity at year-end

8 GW

Up from 3 GW at the end of 2025

China, share of data center revenue

Under 1%


Almost nothing today — and none assumed in the forecast

Four things worth repeating

- **AI stopped being something people type into.** As of last month most AI work is done by agents rather than prompted by a person — and an agent uses 15 to 100 times the computing power of a human doing the same task.
- **Open or closed, NVIDIA gets paid.** It expects both kinds of model to succeed, and says it is the only platform that runs every frontier model.
- **It now gets paid twice.** On some deals NVIDIA sells the hardware and then takes a share of the rental income that hardware earns.
- **It is helping finance its own customers.** About a quarter of next year's business comes from AI labs NVIDIA backs with its balance sheet. Its answer to the circularity worry: the chips can be redeployed to someone else.

Source: NVIDIA Q2 FY27 earnings call, Aug. 26, 2026, as summarized by Jefferies. Management commentary paraphrased, not verbatim.

NVIDIA's Broad Investment Ecosystem

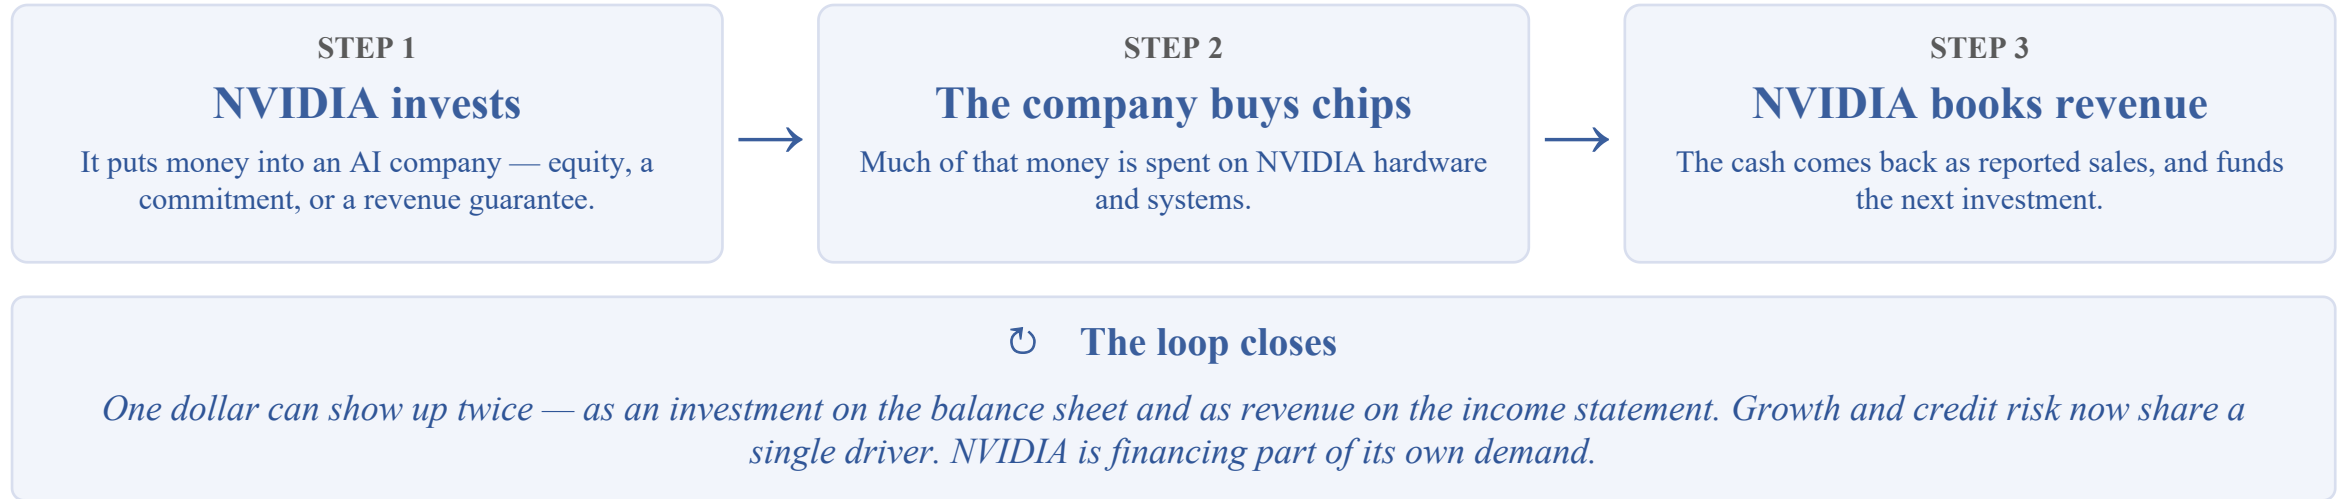


- **Five buckets.** Public equities, frontier labs, supply chain, infrastructure, venture.
- **\$70B+ committed**, about \$18B of it not yet deployed.
- **Public book, \$55B+.** Intel and SpaceX are ~80% of it. Arm fully exited, and the old “90% CoreWeave” story is stale.
- **Frontier labs, ~\$45B.** \$30B OpenAI, closed in March. Up to \$10B Anthropic, up to \$5B Safe Superintelligence.
- **Supply chain.** MediaTek, Corning, Marvell, Lumentum, Coherent, IREN — capital that extends NVIDIA’s standards.
- **Why it belongs here.** NVIDIA is putting money into the companies that buy its chips. When the seller helps fund the buyer, the money can travel in a circle — which is what the next slide unpacks.

Source: Compiled from disclosed positions, public reports and company announcements, as of Aug. 31, 2026. Items marked reported are not officially announced.

What Is Circular Financing?

When the seller helps fund the buyer — and why that makes people nervous



Why it worries people

- Demand can look stronger than it is when part of it is funded by the seller.
- If the AI companies stumble, NVIDIA loses twice — the investment and the sales.
- From the outside it is hard to tell how much of the growth is self-funded.


What NVIDIA says

- About a quarter of next year's business involves labs it backs. The rest does not.
- The chips are fungible — if one customer fails, the compute is redeployed to another.
- It ships to investment-grade customers, or ones backed by an investment-grade partner.

Source: NVIDIA Q2 FY27 earnings call, Aug. 26, 2026, as summarized by Jefferies; press reports. Simplified for illustration — individual arrangements differ in structure.

Anthropic Heads for the Public Market

A \$2 trillion ask, roughly double the valuation of its last private round three months ago



Valuation, \$ billions. IPO figure is a reported target, not a priced deal.

Where things stand

- **A different business model.** NVIDIA can give a model away because it sells the compute. Anthropic sells the model — and is about to ask the public market to price that.
- **The filing is already in.** It submitted a confidential draft S-1 on June 1; the public disclosure documents are expected within weeks.
- **Timing.** That could put the listing as soon as September or early October — inside this quarter.
- **Size.** Up to \$100B raised at roughly \$2T, which would top SpaceX on both counts (\$86B at \$1.77T).

\$11.6B


Second-quarter revenue

More than doubled

Source: The Wall Street Journal (Aug. 25–26, 2026); NPR (June 1, 2026) on the confidential S-1. Figures are reported plans and may change.

The Ultimate Total Addressable Market

A \$30 trillion claim, against \$2.4 trillion of actual revenue for the entire S&P 1500 tech sector



\$ trillions. TAM = annual revenue at a theoretical 100% market share.

How to read a TAM

- **What it is.** The revenue a company could capture at 100% share — a ceiling, not a forecast, built from industry data and bankers’ models.
- **Anthropic’s method.** It is sizing the full scope of work that could be done by AI models. That framing makes the labor market itself the market.
- **The precedent.** SpaceX called its \$28.5T the largest actionable market in human history; \$26.5T of it was AI. Uber claimed \$6T, WeWork \$3T.
- **The pushback.** Damodaran, on SpaceX’s AI number: “reaching the end of what’s plausible and pushing beyond.”

Source: The Wall Street Journal; S&P 1500 tech revenue per FactSet (191 companies, 2025). TAM figures are company claims from IPO materials, not results.

A Live Case for the Fall Program

If the timing holds, Anthropic prices right as the program begins — and trades through it

Program dates

Oct 9 – Nov 20

Eight live 90-minute sessions

Expected listing

Sept – early Oct

Public disclosure documents due within weeks

Coverage universe

~30 companies

China tech and the Magnificent 7

What we expect to do with it

- **Read the disclosure documents as they land.** A live S-1 during the sessions is the best possible teaching text — the first audited look at revenue, gross margin and compute cost at a frontier lab.
- **Test the TAM against the price.** Work backwards from a \$2T valuation to the share, margin and time horizon it implies, and judge whether the \$30T ceiling does any real work in the answer.
- **Watch the aftermarket.** SpaceX jumped, then gave it all back to its \$135 offering price. Let's see whether Anthropic prices well and then trades well — we will follow it together, week by week, through the sessions.
- **China comparison.** The open-weight Chinese labs and the U.S. proprietary labs are pursuing similar demand with different business models. Anthropic's filing is the first time one of them shows its numbers.