Ed Faritov, CFA

Country	United States
Sector	Semiconductors
Issuer	NVIDIA NVDA
Price 6 Nov 2026	$186.64
12m target	$231.04
Financial signal	BUY signal
12m total return	24.32%
Market cap	$4,509.9bn
Dividend yield	0.54% assumed
EV/EBITDA FY26/27E	32.8× / 16.1×
P/E FY26/27E	38.1× / 18.6×
Adjusted net debt/EBITDA	-0.17× FY27E

Source: Valuation and regular-session reference price as of 6 November 2026, 16:00 EST. The cover uses fresh forecast diluted claims and the dated cash/non-operating bridge; these are estimates, not observed daily balance-sheet quantities. The twelve-month horizon is 6 November 2027.

Source: EV is enterprise value; EBITDA is earnings before interest tax depreciation and amortization. P/E is price divided by earnings per share.

NVIDIA business and competitive outlook

Growth is visible; duration needs proof

At the $186.64 close on 6 November 2026, expected intrinsic value is $206.07 and the conditional twelve-month value is $231.04 for 6 November 2027. Including an assumed $1 dividend, the 24.32% total return exceeds the 12.60% required return plus ten percentage points, giving a financial BUY signal. Three reasons govern the stance: segment mix does not establish lasting exceptional rent; customer computing savings do not automatically become NVIDIA profit; and gross buybacks do not establish the continuing owner's economic return. All four business-quality tests remain unestablished; the financial signal does not constitute an active investment recommendation.

More Compute & Networking revenue cannot alone prove a permanent 58% operating margin. Historical segment economics at a 95% share imply about 45.89%; the separate 58% hypothesis adds $39.17 per share to base value. China makes the duration question concrete: models, domestic accelerators, software and clouds are developing a competing route to useful computing. Selected deployments are established; global cost parity and the amount of NVIDIA business displaced are not. Section 5 distinguishes the working links from supply constraints and roadmaps.

Rubin's advantage is conditional on the useful workload. One matched comparison gives about 1.90 times total tokens per dollar and a roughly 53% relative-utilization break-even, using modelled ownership costs. It does not establish universal annual replacement or the selling price NVIDIA captures. A separate one-time 5% capture reduction from FY2028, with retention fixed, costs $9.60 per share of base value.

The owner must also pay for compensation and reinvestment. FY2022–FY2026 gross distributions were 40.69% of profit, but estimated allocation after offsetting new shares was 22.31%; the one-third test remains unestablished. Expected value is $206.07 against the $186.64 close. The FY2029–FY2031 demand-growth adjustment remains 0 percentage points. Holding other inputs and outcome weights fixed, financial BUY requires more than -0.73 points and SELL less than -9.82 points; quality and realization conditions remain separate.

Contents

1. How NVIDIA earns....3
2. Reporting segments and end markets....5
3. The technology and the cost of a task....7
4. What drives the markets....9
5. Can NVIDIA defend its profits....11
6. Supply and market access....13
7. What must be true for shareholders to earn more....14
8. Cash reaches the owner after dilution and reinvestment....15
9. The forecast already requires exceptional expansion....17
10. Growth is visible; price leaves little room for error....18
11. The next evidence must reach deployment and cash....21
12. What can turn the operating success into capital loss....21
13. Sources....22

1. How NVIDIA earns

NVIDIA earns its return by making a customer's computing system useful as a whole. Its graphics processing units, or GPUs, carry out many calculations in parallel. Accelerated computing assigns suitable work to these processors while a general-purpose central processing unit, or CPU, coordinates the system. The sale combines processors, memory connections, networking and software. A customer buying capacity to train or run an artificial-intelligence model ultimately pays for useful work completed at the required quality and speed.

Reporting segment revenue and operating profit are reconciled in Table 12 in the appendix.

The company serves several kinds of buyer. The largest cloud operators build infrastructure that they use internally and rent to other customers. Specialist AI clouds sell access to computing capacity. Industrial and enterprise customers use systems for their own applications. These buyers differ in purchasing power, financing, technical skills and ability to develop an alternative processor. Their investment budgets therefore reach NVIDIA through different routes. A building, an electricity connection and a custom chip can all sit inside a cloud investment budget without generating a dollar of NVIDIA revenue.

Who earns the AI infrastructure dollar

The physical chain begins with a design and ends with a usable application. NVIDIA designs the processor and system and supplies the software platform. Contracted manufacturers fabricate the logic dies; memory suppliers provide high-bandwidth memory (HBM) close to the processor; advanced packaging, including TSMC's CoWoS technology, brings logic and memory together. Systems, interconnects and networking turn packaged processors into an installation. The cloud operator or enterprise must then provide power, cooling, software operation and paying demand. Table 1 shows these relationships. The purchasing order runs back from the buyer, while chips and usable computation move forward.

Table 1: NVIDIA consolidated operating profit was $130.4bn in FY2026; external profit pools remain unmeasured.

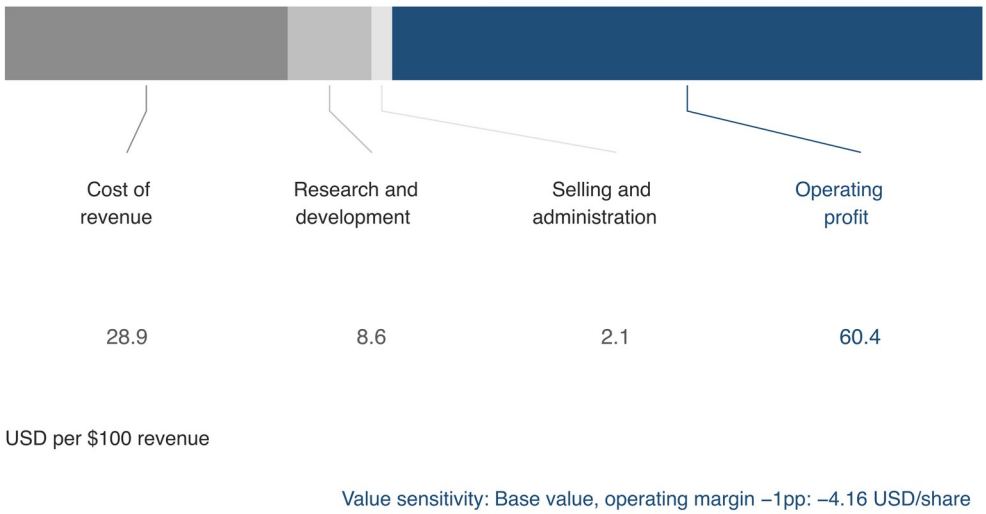
Participant	Economic function	Physical input and output	Measured profit boundary
NVIDIA design and software	Designs accelerated processors systems interconnects and the software platform	Workloads → chip/system designs and software	FY2026 consolidated operating profit $130.4bn; counted here once
TSMC and the contracted wafer manufacturing chain	Fabricates logic dies from the chip design	Chip designs → fabricated logic dies	Comparable NVIDIA-specific profit not measured
High-bandwidth memory suppliers	Provide memory stacks that keep data close to the processor	Memory designs and capacity → HBM	Comparable NVIDIA-specific profit not measured
Advanced packaging chain including TSMC CoWoS	Combines logic dies and memory into the packaged processor	Logic dies + HBM → packaged accelerators	Comparable NVIDIA-specific profit not measured
NVIDIA systems networking and manufacturing ecosystem	Combines accelerators networking and software into a usable computing system	Packaged processors and designs → servers racks and platform software	Included in consolidated NVIDIA; no second allocation
Cloud operators and enterprises	Install and operate systems to supply or consume useful computing capacity	Systems and powered facilities → useful computing	Comparable NVIDIA-specific profit not measured
Application developers and end users	Turn available computation into model outputs and economically useful services	Computing models and demand → applications	Comparable NVIDIA-specific profit not measured

Source: NVIDIA annual and quarterly business disclosures; TSMC Q3 2026 transcript; buyer primary disclosures and original workload study listed in Section 13. External profit pools and NVIDIA allocations are unmeasured, rather than zero.

The profit measurement has a narrower boundary than the physical chain. NVIDIA disclosed FY2026 gross profit of $153.5bn and operating profit of $130.4bn. The latter is after research, development, selling and administration. Comparable NVIDIA-specific profit pools at the foundry, memory, packaging, cloud and application stages have not been measured here. Assigning a percentage of NVIDIA revenue to those external participants would confuse its supplier costs with their profits. The company's software and systems stages are counted once in the consolidated profit pool.

Figure 1 shows why the business attracts capital. In FY2026, each $100 of revenue left $60.4 of operating profit after $28.9 of cost of revenue, $8.6 of research and development and $2.1 of selling and administration. This is the issuer's financial allocation. It shows the scale of rent available to shareholders if it persists; it does not establish which outside participant is the marginal producer or how much rent that participant earns. A one-percentage-point reduction in the future operating-margin path reduces base intrinsic value by $4.16 per share. Figures abbreviate percentage points as pp.

Figure 1: NVIDIA retained $60.4 of operating profit per $100 of FY2026 revenue. Operating profit disappears at total operating costs of $100 per $100 revenue. Base value, operating margin −1pp: −4.16 USD/share.



Source: NVIDIA FY2026 10-K, consolidated income statement

The software connection is central to the sale. CUDA is NVIDIA's programming platform and library environment for GPU workloads. Customers accumulate code, engineering routines and operational experience around it. Changing the processor can require changing that software and retesting its output, as well as replacing physical equipment. These costs can protect NVIDIA's economics, but a successful alternative with sufficient savings can pay for the migration. Switching cost is an economic hurdle whose durability needs evidence.

NVIDIA is treated here as a cyclical technology investment. Large infrastructure orders can accelerate while customers add capacity, then slow when installations, funding or useful demand lag. Durable software advantages can coexist with a sharp equipment cycle. The control points are therefore deliveries, customer deployment, the price of useful computation, cash collection and the margin retained after competitors and suppliers respond. A high recent return does not turn all future expansion into recurring income.

The investment story has four parts. NVIDIA converts a coordinated computing platform into unusually high operating earnings. Expansion should raise earnings per share even as the extraordinary margin gradually recedes and reinvestment consumes cash. The 6 November entry price lies below the probability-weighted value of that demanding path, without establishing lasting rent. Execution and cash collection must sustain

the path; loss of accessible demand or useful-computing advantage would require revised forecasts and a full valuation.

2. Reporting segments and end markets

Two reporting segments and several market platforms describe different views of the same company. Compute & Networking and Graphics are the operating segments used in the financial notes. Market platforms group sales by their use or customer category. This distinction matters because segment operating profit, platform revenue and consolidated profit cannot be interchanged when assessing margins or building a forecast.

Table 2 retains the two operating segments and shows the corporate reconciliation. Compute & Networking generated $193.5bn of FY2026 revenue and $130.1bn of segment operating profit. Graphics generated $22.5bn and $9.2bn respectively. Their combined segment operating profit of $139.3bn exceeded consolidated operating profit by $8.9bn. Those central costs matter to the owner even though the segment presentation places them outside the two results.

Table 2: Two reporting segments account for $215.9bn of FY2026 revenue before separate corporate profit reconciliation.

Period	Segment or reconciliation	Revenue $bn	Operating profit $bn
FY2026	Compute & Networking	193.5	130.1
FY2026	Graphics	22.5	9.2
FY2026	Corporate reconciliation	—	-8.9
FY2026	**Consolidated NVIDIA**	**215.9**	**130.4**
Q2 FY2027	Compute & Networking	88.3	62.7
Q2 FY2027	Graphics	7.9	3.9
Q2 FY2027	Corporate reconciliation	—	-2.9
Q2 FY2027	**Consolidated NVIDIA**	**96.2**	**63.7**
H1 FY2027	Compute & Networking	162.9	116.0
H1 FY2027	Graphics	15.0	6.8
H1 FY2027	Corporate reconciliation	—	-5.6
H1 FY2027	**Consolidated NVIDIA**	**177.8**	**117.3**

Source: NVIDIA FY2026 Form 10-K and Q2 FY2027 Form 10-Q; reported USD bn, with corporate reconciliation calculated from the disclosed totals.

Data Center generated $164.3bn of H1 FY2027 revenue, comprising $91.8bn from Hyperscale and $72.5bn from AI Clouds, Industrial and Enterprise; Edge Computing contributed a further $13.6bn. Table 3 places these current market-platform categories alongside the historical platforms. Hyperscale identifies the largest cloud customers, while AI Clouds, Industrial and Enterprise, abbreviated ACIE, groups other AI-cloud and enterprise demand. Together the current Data Center and Edge platforms reconcile to reported H1 revenue of $177.8bn. They are not additional operating segments.

Table 3: Current Data Center revenue was $164.3bn in H1 FY2027; historical market platforms use a different perimeter.

Period	Market platform	Revenue $bn
FY2026	Data Center	193.7
FY2026	Compute (within Data Center)	162.4

Period	Market platform	Revenue $bn
FY2026	Networking (within Data Center)	31.4
FY2026	Gaming	16.0
FY2026	Professional Visualization	3.2
FY2026	Automotive	2.3
FY2026	OEM and Other	0.6
Q2 FY2027	Data Center	89.0
Q2 FY2027	Hyperscale (within Data Center)	48.7
Q2 FY2027	AI Clouds Industrial and Enterprise (within Data Center)	40.3
Q2 FY2027	Edge Computing	7.2
H1 FY2027	Data Center	164.3
H1 FY2027	Hyperscale (within Data Center)	91.8
H1 FY2027	AI Clouds Industrial and Enterprise (within Data Center)	72.5
H1 FY2027	Edge Computing	13.6

Source: NVIDIA FY2026 Form 10-K and Q2 FY2027 Form 10-Q; reported USD bn, with corporate reconciliation calculated from the disclosed totals. Parent and included rows are not added together; the Q2 classification is used for current platforms.

The historical platform view remains useful for understanding the franchises. In FY2026, Data Center revenue was $193.7bn, including $162.4bn of compute and $31.4bn of networking. Gaming contributed $16.0bn, Professional Visualization $3.2bn, Automotive $2.3bn and OEM (original equipment manufacturer) and Other $0.6bn. Gaming serves the consumer graphics and computer market; professional visualization serves workstations and related applications; automotive includes computing platforms used in vehicle and driving applications. These activities have different product cycles, but their old lines are not a disclosed current FY2027 breakdown.

The FY2027 presentation combines Data Center, split into Hyperscale and ACIE, with Edge Computing. The Q2 disclosure also moved one customer from ACIE to Hyperscale and recast comparisons. Quarter-to-quarter analysis therefore uses that Q2 classification. Joining an originally published Q1 platform split to the recast Q2 series would create a spurious growth signal. The current verified disclosure also does not supply the old compute/networking split, so no assumed networking attachment ratio is presented as reported fact.

This reporting change affects how investors read the forecast. The Hyperscale and ACIE paths in Section 9 represent customer-market exposures with different substitution and deployment risks. They do not provide separately disclosed future operating margins. A larger share of the company in Compute & Networking may support consolidated profitability, but the amount depends on segment economics and central costs. Section 5 measures that distinction; simply calling more revenue “Data Center” cannot establish a permanent high margin.

Customer concentration also crosses the reporting boundaries. Three direct customers represented 16%, 15% and 13% of H1 revenue. Five direct customers represented 22%, 14%, 13%, 11% and 10% of accounts receivable. Direct customers can include distributors or intermediaries, so these figures do not precisely identify end users. They nevertheless show why execution and collection at a small set of counterparties can materially affect the quarter. The separately priced fifteen-day receivables delay reduces base value by $2.53 per share.

3. The technology and the cost of a task

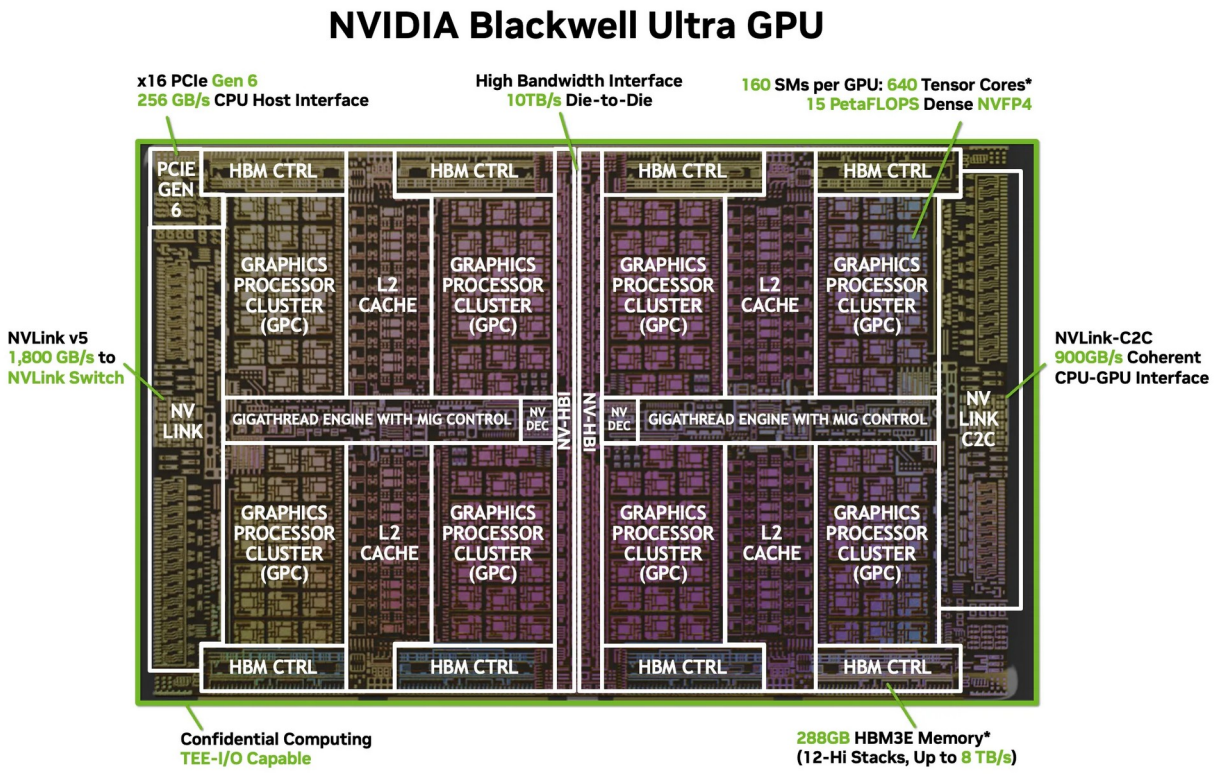
The economic unit is a useful result delivered under a stated service requirement. Training adjusts a model using examples; inference runs the trained model to produce an answer or decision. A system that completes more arithmetic may still deliver less economic value if it runs out of memory, waits for data, misses the required response speed or remains idle. This is why product comparisons must retain their workload, accuracy, system size and operating assumptions.

System performance and product economics

A memory controller manages movement between HBM and the computing blocks. The package and its connections determine whether those blocks can obtain data fast enough to do useful work. NVLink is NVIDIA's high-speed connection between processors and systems. These components address different bottlenecks: adding computing units does little if the application spends its time waiting for data or for another processor.

Figure 2 is NVIDIA's own Blackwell Ultra architecture illustration. It explains how two compute dies, memory controllers and system links fit together. The diagram's headline specifications describe the manufacturer's architecture; they are not an independent measurement of application throughput, a customer invoice or a margin forecast. For the investor, the relevant connection is that performance increasingly depends on the package and the surrounding system, creating more components that must work together and more interfaces that a rival must match.

Figure 2: Blackwell Ultra joins 2 compute dies, memory controllers and system links on one GPU package. Manufacturer architecture illustration; no standalone valuation input.



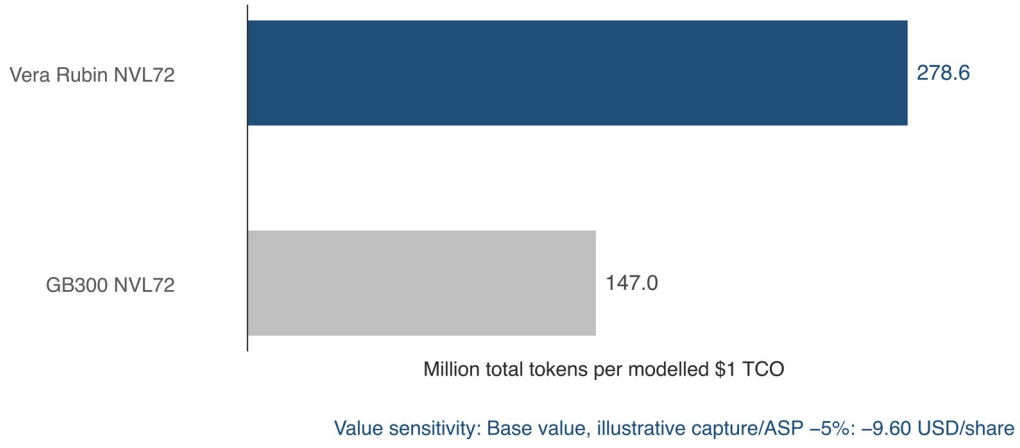
Source: Original source; original title “Figure 1. NVIDIA Blackwell Ultra GPU chip explained”, 22 August 2025, updated 24 September 2025. This diagram is not a workload benchmark.

The smaller labels in the illustration describe that machinery. PCI Express (PCIe) connects the GPU to the host CPU; NVLink-C2C is a chip-to-chip link. A graphics processor cluster (GPC) contains streaming multiprocessors (SMs), which schedule and execute computing work. Tensor Cores specialize in matrix arithmetic. L2 is on-chip cache, which reuses nearby data. The GigaThread engine schedules work; MIG, or Multi-Instance GPU, partitions a GPU into isolated resources. NVDEC and NVENC decode and encode video. TEE-I/O denotes protected input and output for a trusted execution environment. HBM3E identifies the memory generation, while stacked memory layers increase capacity. Gigabytes per second (GB/s) and terabytes per second (TB/s) measure data bandwidth; floating-point operations per second, or FLOPS, measure arithmetic capacity. NVFP4 and FP8 use four-bit and eight-bit numerical formats respectively, so any speed comparison must retain the required accuracy.

An older “30 times” comparison measured a complete Blackwell system with 72 GPUs against a Hopper system with eight. In the same MLPerf Inference v5.0 work, the Llama 3.1 405B result was up to 3.4 times per GPU in server mode and 2.8 times offline. The system result therefore does not establish a thirtyfold gain for each processor or each dollar spent. The Rubin evidence used here is a next-generation system preview published on 19 October 2026; a preview comparison does not establish the customer's realized purchase price, installation date or production utilization.

Figure 3 shows the closest current comparison in the available evidence: one matched DeepSeek-V4-Pro workload from the original InferenceX study. At 100 tokens per second per user, Rubin NVL72 delivered 278.6 million total tokens per dollar against 147.0 million for GB300 NVL72, a ratio of approximately 1.90. A token is a unit of text handled by the model. About 99% of the counted tokens in this test were input or cached input, so the result is not the cost of an output token. The serving software also differs between systems. ASP in the sensitivity label means average selling price; its change is a separate hypothetical pricing assumption.

Figure 3: Rubin yields 278.6 versus 147.0 million total tokens per dollar on one matched workload. Rubin advantage disappears below approximately 53% relative utilization, holding the stated TCO and workload fixed. Base value, illustrative capture/ASP −5%: −9.50 USD/share.



Source: Original source

Total cost of ownership, or TCO, combines the assumed cost of equipment and operation. The study models $3.61 per GPU-hour for Rubin against $2.31 for GB300. Their ratio is about 1.56. Combining the published rounded cost and token figures implies roughly 2.96 times useful throughput per GPU-hour in the selected workload. At those costs and the same quality and response requirement, Rubin needs more than 1.56 times throughput to break even. Equivalently, its relative useful utilization must remain around 53% or higher under the study's linear-utilization assumption.

The service-level requirement can reverse the result. In the reported DeepSeek curve, GB300 wins at a much tighter interactivity requirement, approximately above 256 tokens per second per user. The GLM, Qwen and MiniMax comparisons are excluded from the matched economic claim because they involve different model checkpoints, broad interpolation or an older comparator run. A checkpoint is a particular saved version of a trained model. Restricting the claim to a comparable task gives investors a usable condition; it does not establish which system wins every task.

NVIDIA is fabless: it designs chips and contracts manufacturing rather than owning the wafer plants that produce them. Its supply commitments give it contractual exposure to manufacturing, packaging and memory capacity, while actual available allocation to NVIDIA remains unmeasured in this analysis. An application-specific integrated circuit, or ASIC, is designed for an intended class of workloads. Google's tensor processing unit, or TPU, and Amazon's Trainium are relevant alternatives; their economics depend on the task, software and deployment scale rather than a universal cost advantage.

The customer's saving and NVIDIA's profit require separate evidence. A faster system can support a higher price, enable a customer to buy fewer systems, attract new workloads or lose some benefit through installation and idle time. The evidence does not measure the price elasticity that divides those effects. Our independent one-time 5% reduction in NVIDIA's realized capture from FY2028 lowers base value by $9.60 per share. That is an explicitly chosen pricing and mix stress, with ASIC retention held fixed; it is not an elasticity estimated from the 1.90-times benchmark.

A Chinese AI model is not necessarily running on a Chinese chip. DeepSeek's V4.1-Flash reduces the memory needed to retain context: at equal sequence lengths, its global cache is about one quarter of its predecessor's and persistent storage about one eighth. Those are specific cache savings, not a fourfold reduction in all hardware. Its October DeepInfer software supports both NVIDIA SM100 and Huawei Ascend 950 processors. A more efficient model can reduce computing per task, encourage additional tasks, or do both; the response of total paid demand has not been measured. Model progress therefore challenges assumptions about hardware demand without establishing NVIDIA displacement. DeepSeek report, pp. 1, 4–5; DeepInfer.

4. What drives the markets

Buyer investment is expanding rapidly, but the economic route from that investment to NVIDIA remains selective. Four large US buyers acquired $201.9bn of property and equipment in the quarter ended September 2026, up from $97.3bn a year earlier. Their combined operating cash flow was $175.5bn. The $26.4bn difference shows that part of their gross cash investment needed another source of funding; it is not a consistently defined combined free-cash-flow measure or a GPU purchasing budget.

Capacity spending and future supply

The accounting perimeter matters. Microsoft reported $39.4bn of cash property-and-equipment purchases, Alphabet $51.4bn, Amazon $77.3bn and Meta $33.9bn. These cash acquisitions exclude new lease recognition and lease-principal repayments for comparability. Meta's separately stated capital expenditure includes $1.0bn of finance-lease principal beyond that cash-PPE figure. Amazon's gross figure includes businesses outside Amazon Web Services (AWS). Adding different headline capex definitions would exaggerate comparability before any attempt to allocate the spending to processors.

The contents of the budget are equally important. Microsoft described roughly three-quarters of its quarterly investment as short-lived assets, predominantly processors, memory and storage. That percentage is not NVIDIA's share. Alphabet described approximately 65% of technical-infrastructure investment as

servers and networking and 35% as data centres. Its denominator is technical infrastructure rather than every company asset, and the servers include its own accelerators. The appropriate revenue question is how much useful computing demand remains accessible to outside suppliers and how much NVIDIA wins within it.

A capacity cycle has several dates. A buyer can reserve power or a building, order equipment, receive a system, connect it, fill it with workloads and finally collect revenue from customers. The cash dates need not coincide. Strong orders can therefore precede an interval in which customer financing and deployment become more important than the speed of the chip. The next evidence should establish which stage has advanced. Treating every announced infrastructure dollar as an immediate NVIDIA order skips that economic sequence.

Competition is broadening the accessible market's boundary. Alphabet reported about one hundred external TPU customers and plans to expand sales. Amazon intends to offer Trainium beyond AWS in late 2027. This makes external custom-accelerator supply a relevant competitive route. It does not establish how many NVIDIA sales disappear: those systems may serve another workload, replace general-purpose computing or make a previously uneconomic application possible. Revenue forecasts therefore separate growth in buyer computing demand, the part retained by the merchant market and NVIDIA's capture of that market. Merchant means equipment sold to outside customers rather than kept entirely inside its developer.

Chinese clouds add another route: the equipment buyer can develop its own processor and sell the resulting computing service. Alibaba reported RMB12.4bn of AI-related product revenue in the June 2026 quarter and describes orchestration across different chips, including its own. Its Zhenwu processors serve external customers through Alibaba Cloud. This creates an incentive to retain more computing revenue inside the cloud rather than pass it to an outside chip supplier. It does not establish how much NVIDIA equipment was replaced: the service can also meet additional demand. The forecast must distinguish market growth from the share still accessible to NVIDIA. Alibaba results, pp. 3–4.

The two large Data Center categories deserve different assumptions. Hyperscale buyers have engineering scale and an incentive to diversify supply or integrate their own designs. ACIE contains specialist clouds and enterprises with a different mix of skills, funding and deployment timelines. The base forecast reaches FY2028 revenue of $698.1bn using the preliminary company growth anchor; subsequent paths incorporate separate retention and growth assumptions. A 5% reduction in the Hyperscale retained-demand multiplier during FY2028–FY2031, with capture unchanged, reduces base value by $16.57 per share.

AI adoption and the limits of the growth history

The available history does not establish a completed AI-infrastructure life cycle. FY2022–FY2026 contains a weak year and a dramatic expansion, but it does not measure ultimate application penetration, the settled replacement interval or steady-state competition. The long forecast is therefore a conditional path rather than a claim that adoption has been measured to maturity. The 3% nominal growth used after the explicit horizon is a valuation convention. Raising that convention to 3.5% adds $3.71 per share in a separate complete revaluation; that sensitivity does not quantify the cost of not knowing the complete cycle.

The present investment stance follows this separation of dates and markets. Preliminary guidance supplies a near-term anchor, buyer spending demonstrates physical commitment, and later accessible demand remains a hypothesis that must survive deployment and competition. A five-percentage-point reduction in annual Hyperscale and Industrial demand growth during FY2029–FY2031 reduces base value by $19.42 per share. The investment decision changes through those specific revenue and cash consequences, rather than through a general label that the AI market is strong or weak.

5. Can NVIDIA defend its profits

NVIDIA has converted the infrastructure expansion into exceptional reported profits. FY2026 revenue reached $215.9bn and operating profit $130.4bn, compared with $27.0bn and $4.2bn in FY2023. H1 FY2027 then produced $177.8bn of revenue and $117.3bn of operating profit. These results demonstrate present scale and operating leverage. The harder investment question is how much of the improved economics survives after customers have more computing capacity and credible alternatives.

Segment mix and the durability of margins

A larger Compute & Networking share alone does not establish a higher lasting margin. The reporting-segment history allows a useful challenge to the idea that a larger Compute & Networking share alone explains a lasting 58% consolidated operating margin. Applying historical segment margins and central costs to a 95% Compute & Networking revenue share produces approximately 45.89%. Even a 100% share produces about 46.59%. A 58% result would require an impossible share of about 181% if those historical relationships were held fixed. This rejects that particular mix explanation; it does not prove that NVIDIA cannot earn 58% under a genuinely different economic regime.

The strong alternative is a durable improvement within the segment itself. Better software integration, system value, lower relative central costs or a longer period of constrained competitive supply could support more rent than the historical relationship implies. That is a separate proposition investors must pay for and monitor. Using 58% as the conditional long-run margin raises current base value from $209.33 to $248.51, a $39.17 difference. We would change the historical-margin judgement if reported segment economics and central-cost absorption demonstrate a durable improvement as capacity and competitors expand.

Data Center carries the largest economic exposure. Strategy is to make each generation's useful computing improvement available through an integrated platform: processors, memory connections, networking, system design and software. A faster design must arrive in a deployable system, and contracted supply must support delivery. The competitive defence is strongest where replacing that system entails software migration and operational risk. The strategy therefore depends on customer economics and reliable execution; an unmeasured majority of manufacturing capacity or an assumed universal annual replacement cycle cannot establish its success.

China's challenge is a complete computing chain

China's competitive ambition runs from models through chip design, manufacturing, systems and software to the application customer. Each link matters because the customer buys completed work at an acceptable cost and speed. A capable model alone cannot deliver that service; a capable accelerator without usable software, memory or reliable supply cannot deliver it either. The emerging chain gives NVIDIA's customers another option, while its incomplete links limit how quickly that option can scale.

The evidence is already stronger than a laboratory announcement. Huawei and SiliconFlow describe serving DeepSeek-R1 on a CloudMatrix system; Alibaba supplies computing based on its own designs. Open models and portable software make these systems easier to use. However, these observations establish particular workloads and commercial routes, rather than universal equivalence to NVIDIA's newest platform.

Table 4: Six links make a Chinese alternative possible; capacity and comparable economics still constrain substitution.

Link	What already works	What limits the conclusion	Meaning for NVIDIA
Models	DeepSeek and Qwen offer downloadable models; Qwen3.8-27B has an Apache 2.0 licence.	Licensing varies across models; weights still require computing and support. Benchmarks do not measure procurement decisions.	More model choice can expand GPU workloads or make alternatives easier to adopt. Qwen model card.
Chip design	Alibaba's Zhenwu computing is commercially offered through its cloud.	A design and customer count do not establish useful capacity, production cost or displaced NVIDIA orders.	A large cloud can keep more of the computing budget inside its own stack.
Wafers, memory and packaging	SMIC manufactures semiconductors; JCET reports volume production of advanced packaging technology.	SMIC's aggregate capacity is not AI-chip yield or allocation. JCET does not identify Ascend packaging; available memory stacks and finished AI packages are unquantified.	Physical supply limits the pace of substitution. SMIC interim report, pp. 7, 13–14; JCET interim report, p. 12.
Systems and networks	CloudMatrix runs large-model inference by coordinating processors, memory and networking.	Supplier-authored tests use particular workloads and configurations; they do not establish comparable Blackwell or Rubin ownership costs.	A competitive system can substitute for selected tasks without winning every chip-level comparison. CloudMatrix study, §§3.3, 5.1–5.3.
Software	DeepInfer supports both NVIDIA and Ascend hardware.	Published performance plots use NVIDIA; the open release omits components of a complete service. Migration cost and reliability remain workload-specific.	Easier porting can reduce switching costs; NVIDIA remains a supported destination.
Cloud and applications	Alibaba's cloud operates in 31 regions and offers tools for international model-service providers.	Region count is not paid AI usage or the accelerator mix. Its V900 commercial release remains a Q1 2027 plan.	Competition can reach foreign customers through a service without a direct chip sale. Alibaba international update.

Source: The primary model, semiconductor, system and cloud disclosures linked in the rows and surrounding discussion. This comparison records disclosed products, tests and plans; it does not allocate revenue, profit or market share.

NVIDIA's protection consequently rests on the cost of switching a production workload: software coverage, engineering time, useful utilization, reliable delivery and support. Chinese alternatives weaken that protection where a customer can repeat a large, well-defined task economically. They need not reproduce every CUDA capability to win such a task. Conversely, a single optimized demonstration does not show that customers can move a diverse fleet cheaply. The decisive comparison includes accepted output quality, latency, migration, downtime and power, alongside equipment cost.

This also changes the buyer's bargaining position before a complete migration. A credible second source can press NVIDIA to share more of its productivity gain through price or service. The relevant question is whether a 10% NVIDIA price increase would cause a customer to move useful work. The evidence does not yet measure that response. Suppliers of suitable wafers, memory and packaging also retain bargaining power: developing an alternative chip does not remove those dependencies.

The valuation separates these channels. Its existing retention sensitivity applies a 0.95 multiplier to Hyperscale retention in each year from FY2028 to FY2031 and costs $16.57 per share of base value; a separate one-time 5% capture reduction from FY2028 costs $9.60; bringing margin normalization forward to FY2032 costs $8.81. These are conditional illustrations, not estimates of Chinese competition, and cannot be added without a joint recalculation. The evidence supports monitoring retention, realized pricing and the duration of exceptional margins. A China-specific loss and its probability have not been measured. These observations do not supply a new forecast coefficient or establish that the conditional valuation includes a quantified China case.

The cadence of product introductions creates an opportunity and a liability. Successful new systems can enlarge useful capacity and accelerate orders. Transition costs, inventory, product qualification and the customer's installation schedule can delay revenue or reduce margins even when a product is technically attractive. The Q3 FY2027 company outlook for $108.0bn of revenue is an execution anchor. Its preliminary FY2028 expectation of approximately 70% revenue growth is stronger near-term evidence than an unsupported extrapolation of an older growth rate, but it remains an expectation.

The other end markets deserve economic context without fabricated precision. Gaming depends on consumer computer demand, product refresh and channel inventory. Professional visualization depends on workstation spending and application adoption. Automotive follows design, qualification and deployment programmes that can take longer to mature. Their historical revenue is in Section 2; the current Edge presentation does not disclose a clean quarterly split among them. Consequently, the forecast treats Edge as the disclosed platform and does not attach an unreported current profit margin to each old franchise.

The worst observed year supplies a partial check on the advantage. NVIDIA's FY2023 gross margin was 56.93%, against AMD's 44.93% for calendar 2022, a spread of about 12.00 percentage points. Adding back AMD's acquisition amortization in cost of revenue reduces the comparison to about 5.87 points. The sign survives that adjustment, but product mix, acquisitions and different year ends prevent it from proving a matched GPU cost advantage. A positive consolidated margin is insufficient evidence for the stricter product-level test.

Competitive durability is therefore measured through retention, segment economics and the length of the high-return period. Bringing margin normalization forward to FY2032 reduces base value by $8.81 per share. Eliminating margin fade adds $61.73. These large differences explain why product leadership and the price of its duration must be considered together. The current price does not justify increasing exposure solely because the largest segment is growing faster than the rest of the company.

6. Supply and market access

The external environment acts through delivery, demand and the cash that reaches owners. Export rules determine which products can reach which customers; manufacturing and packaging determine available systems; financing determines which customers can deploy what they order. These channels can reinforce or offset each other. A favourable chip benchmark cannot compensate for a system that cannot legally ship or a customer unable to finance its installation.

PEST

Political, economic, social and technological forces, together with legal conditions, are the subjects of PEST analysis. China remains a real revenue geography even when a particular Data Center product is excluded from an outlook. In Q2 FY2027, revenue attributed to direct customers headquartered in China and Hong Kong was $7.9bn. H200 sales were less than 1% of Data Center revenue, and a $0.4bn related charge had

already occurred in H1. The company's Q3 outlook assumes no Data Center compute revenue from China. That is a product-specific forecasting boundary, not a statement that all Chinese revenue is zero or that future licensing has a known value.

Physical and legal access constrain this competitive route. Suitable compute dies, memory, packaging, networks and power must arrive together; general semiconductor capacity does not establish available AI systems. Foreign availability also needs confirmation: Huawei's announced date for AICS outside China is 30 November 2026, still ahead of this report's information cut-off. Export restrictions can affect both NVIDIA's sales and the alternative supply chain. NVIDIA's Q3 outlook already assumes no Data Center compute revenue from China, so deducting that same excluded Q3 revenue again would count the exclusion twice. Competitive pressure on other revenue or future pricing is a separate question. The additional investment risk is future loss of tasks, pricing power or margin duration, including outside China. Huawei availability announcement; NVIDIA Q2 outlook, 26 August 2026; NVIDIA Q2 Form 10-Q, pp. 27, 37–38.

The geographic presentation also requires care. Q2 revenue attributed to US-headquartered direct customers was $60.1bn and to Taiwan $27.0bn. These addresses identify the direct customer rather than the ultimate location of useful computing. They should not be used to measure an exact end-demand share or to claim that a manufacturing location is the consuming market. Regulation and concentration can still affect collection and delivery through those intermediaries.

The supply chain concentrates several physical risks. Wafer fabrication, advanced packaging, memory and the customer's powered data centre must all arrive on compatible schedules. TSMC's disclosed $60.0–64.0bn capital expenditure plan for 2026 demonstrates investment in supply, but the reviewed primary material does not quantify NVIDIA's absolute CoWoS allocation or its HBM allocation. Capacity announcements therefore do not produce a defensible NVIDIA shipment ceiling here. A supply or geopolitical disruption can affect both the equipment and the customer's ability to use it, while an arrival of abundant capacity can change pricing and the duration of rent.

Commitments make a delivery interruption costly even for a company without wafer plants. NVIDIA disclosed $366.0bn of supply, cloud, lease, equity and capital commitments, plus $56.0bn of additional AI-cloud and lease commitments. They do not share one payment date or cancellation right. Conditional guarantees of $108.5bn are another exposure, rather than current financial debt. The appendix preserves the category reconciliation and separately priced timing and loss assumptions.

The return required by the investor also matters. The 6 November US Treasury ten-year rate of 4.86%, the retained measured beta of 1.72 and the common 4.50% equity-premium convention produce a 12.60% required equity return. Beta is the source 260-week estimate, not a fresh 6 November resample. A higher financing hurdle can absorb operating success before it reaches the share price. Current capital weights, alternative premiums and investment-asset classifications are replayed in the complete dated valuation.

7. What must be true for shareholders to earn more

Strong AI demand is already part of the price. Increasing exposure requires evidence that NVIDIA can retain more useful computing demand or more profit from it than the current forecast assumes. Five operating questions decide whether that evidence is arriving.

Do alternatives take orders away? Google's TPU, Amazon's Trainium and China's developing computing ecosystem can reduce the market open to NVIDIA; migration costs and system performance can protect its position. The scale of substitution remains unmeasured. The existing standalone retention stress reduces base value by $16.57 per share. A lower realized selling price is a different channel: the separate capture

stress costs $9.60. Neither experiment measures a China-specific loss, and the two cannot simply be added. Section 5 examines how Chinese models, processors, software and clouds can make that substitution credible.

Does cheaper computing create enough new demand? A lower cost per useful task can enable new applications and greater consumption, but the evidence here does not establish a universal spending response. The matched workload in Section 3 demonstrates a conditional customer saving. Whether total expenditure rises depends on how many additional useful tasks customers buy, the selling price they pay and the utilization they achieve. Consequently, neither an automatic annual replacement cycle nor a fixed task-cost decline is used as an observed law of revenue growth. The purchase argument fails if the deployed workload loses its saving at the actual service requirement.

Can supply grow without giving away the rent? New manufacturing capacity can improve deliveries and reduce scarcity simultaneously. Shipments alone do not prove pricing power. Section 5 separates revenue mix from the underlying margin advantage.

Can buyers deploy and pay? Orders need power, useful workloads and funding before they become collected cash. The independent fifteen-day collection delay costs $2.53 per share of base value.

Can a growing business still disappoint its owner? Yes: the forecast more than triples FY2026 revenue by FY2028, while margins decline and reinvestment absorbs cash. The action thresholds and next dated test are in Section 10; growth by itself does not cross them.

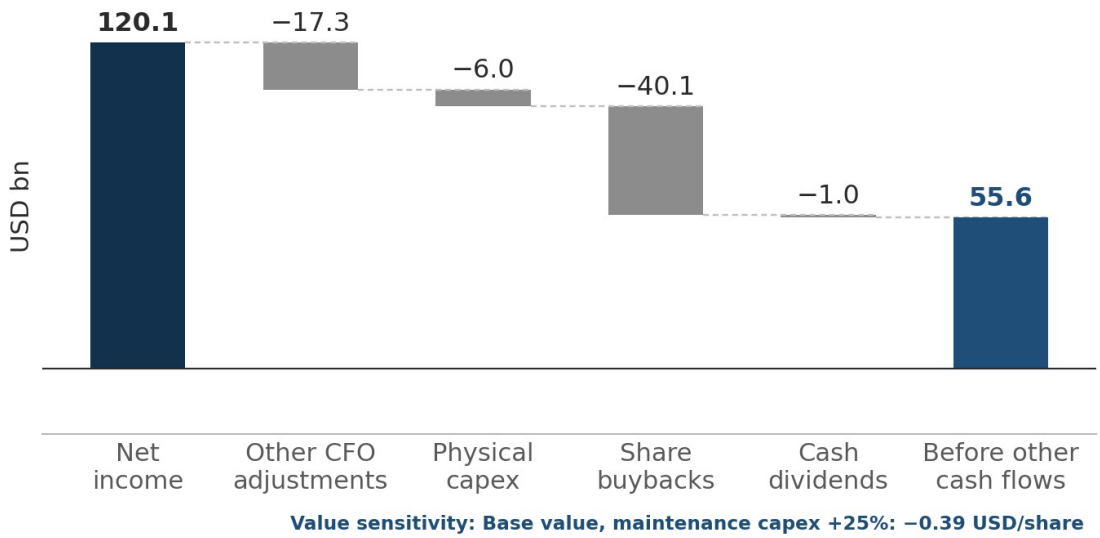
8. Cash reaches the owner after dilution and reinvestment

The shareholder owns what remains after NVIDIA pays to keep earning. Compensation in shares, replacement of existing assets and working capital all use part of that economic return. Gross repurchases therefore need to be read alongside the shares created, rather than treated as an automatic benefit to the continuing owner.

Cash generation and shareholder distributions

Figure 4 follows FY2026 profit into cash. Net income of $120.1bn became $102.7bn of operating cash. Physical investment of $6.0bn left $96.7bn before distributions; repurchases of $40.1bn and dividends of $1.0bn left $55.6bn before other investing and funding. Securities purchases and other transactions remain outside this bridge, so the last figure is not the actual change in cash.

Figure 4: FY2026 left $55.6bn after physical capex and shareholder payouts, before other investing and funding. Cash remaining turns negative if the same-basis payouts exceed $96.7bn. Base value, maintenance capex +25%: −0.39 USD/share.



Value sensitivity: Base value, maintenance capex +25%: −0.39 USD/share

Source: NVIDIA FY2026 10-K, consolidated statements of income and cash flows

The historical payout claim weakens when dilution is included. FY2022–FY2026 gross distributions were 40.69% of profit. Allocating repurchases first to offset estimated new shares leaves 22.31% for dividends and purchases that enlarge the continuing owner's claim; shares outstanding fell only 1.96%. This estimate falls below the one-third quality threshold. It would need another $26.1bn of economic distributions, or a reduction in the estimated dilution-offset allocation from $43.5bn to $17.4bn, to pass on the same earnings. The historical shortfall is not deducted again from today's value.

High recent returns do not close the other quality questions. Five-year average return on invested capital of 113.2% exceeds the current 12.60% hurdle, but a complete AI cycle and historical hurdle comparison are unestablished. The sustaining/growth investment split is undisclosed, and the worst-year margin comparison in Section 5 is only a partial product-economics proxy. All four tests therefore remain unestablished; this limits the endorsement rather than proving a poor business.

Reinvestment and labour costs remain real even in a fabless company. FY2028 physical investment is estimated at $17.0bn, including $5.5bn of maintenance and $11.5bn for growth; research and development of $55.8bn is already expensed. Raising the share-compensation expense by one percentage point reduces base value by $4.16 per share. Detailed asset lives, maintenance alternatives and cash-conversion controls are in the technical appendix.

Collection must now keep pace with earnings. FY2028 assumes sixty receivable days, 115 inventory days and fifty-seven payable days; later inventory days decline toward ninety. A forecasting relationship tested against FY2026 operating cash overpredicts it by $4.1bn. The evidence has not established whether that miss should persist. Its temporary and permanent transfer alternatives are retained in the appendix, rather than presented as a measured probability range.

The assumed $1 annual dividend annualizes the current $0.25 quarterly policy; future declaration and timing are not promised. Buybacks already affect forecast earnings per share. Available owner cash, paid dividends and the benefit of fewer shares must therefore remain separate when judging return.

9. The forecast already requires exceptional expansion

NVIDIA must turn announced demand into a much larger delivered and collected business. The forecast reaches $410.7bn of revenue in FY2027 and $698.1bn in FY2028, then $1,190.2bn in FY2031. Growth on that scale is part of the starting investment case, not an additional reason to pay more for the shares.

H1 FY2027 reported revenue of $177.8bn and Q3 guidance of $108.0bn anchor the immediate path. Q3 guidance has a ±2% revenue range, a 74% gross-margin midpoint and $9.2bn of expenses on the reported accounting basis, including share compensation. The estimated $124.8bn Q4 is the analyst's delivery forecast, rather than disclosed orders. The quarterly bridge, platform growth assumptions and fiscal-calendar reconciliation are retained in the appendix.

The base revenue path reaches $410.7bn in FY2027 and $698.1bn in FY2028. FY2027 combines the reported H1, the Q3 guidance anchor and an explicitly estimated Q4. FY2028 retains management's preliminary expectation of roughly 70% total growth. Subsequent years separate expansion in customer computing demand, the share remaining accessible to outside suppliers and NVIDIA's capture within that share. Hyperscale and ACIE have different paths because their customers' alternatives and deployment constraints differ. These later assumptions are forecasts requiring evidence, rather than observed spending elasticities or an automatic consequence of faster processors.

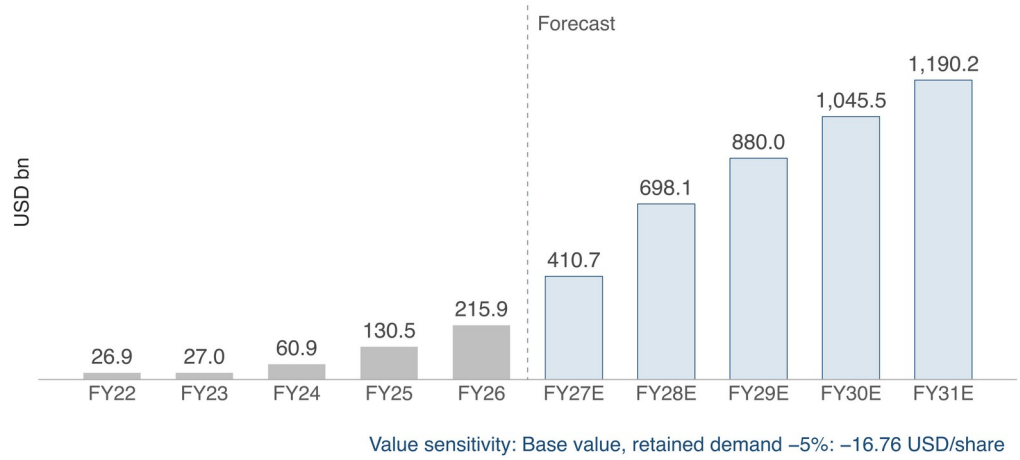
Table 5: Strong growth is already required; operating margin declines as supply and customer alternatives broaden.

Metric	FY2026	FY2027E	FY2028E	FY2031E
Revenue $bn	215.9	410.7	698.1	1,190.2
EBITDA $bn	133.2	271.8	446.9	690.1
EBITDA margin	61.70%	66.18%	64.01%	57.98%
Operating profit $bn	130.4	267.1	439.3	673.3
Operating margin	60.38%	65.05%	62.93%	56.57%
Net income $bn	120.1	242.7	362.1	558.2
Operating cash $bn	102.7	192.3	351.4	584.1
Firm free cash flow $bn	—	—	317.2	513.7
Physical investment $bn	6.0	11.9	17.0	17.9
Financial debt $bn	8.5	33.4	33.4	33.4
Cash and debt securities $bn	49.7	102.8	306.2	1,260.9
Net financial debt / EBITDA	-0.31×	-0.26×	-0.61×	-1.78×
Total dividends $bn	1.0	18.4	24.1	22.7
Fiscal dividend per share $	0.04	0.76	1.00	1.00
Yield at $186.64	0.02%	0.41%	0.54%	0.54%

Source: Operating forecast valued at 6 November 2026 market and financing inputs; E denotes estimates. Firm free cash flow is available to debt and equity after operating reinvestment. Liquidity excludes equity investments. The net-debt measure excludes leases, acquisition consideration and guarantees. FY2027 fiscal dividend $0.76 differs from the forward twelve-month $1 assumption. Dashes mean unavailable; the complete forecast through FY2041 accompanies the calculation.

Five mechanisms govern the path. **Customer mix:** FY2028 comprises $398.0bn of Hyperscale, $267.6bn of AI Clouds, Industrial and Enterprise, and $32.5bn of Edge. Customers have different alternatives, funding and installation constraints. Figure 5 retains the full revenue path; these categories do not carry separately disclosed future margins.

Figure 5: FY2028 revenue reaches $698.1bn on the guidance-anchored path. The later demand and capture assumptions remain conditional. In the current 6 November base case, retained demand −5% reduces V_0 by $16.76 per share.



Source: NVIDIA annual reports and retained operating forecast, with the complete valuation graph dated 6 November 2026.

Accessible demand: more useful computing must translate into spending open to outside suppliers and orders NVIDIA can retain. The standalone five-percentage-point reduction in annual Hyperscale and Industrial demand growth in FY2029–FY2031 costs $19.42 per share of base value. It is a sensitivity to a specified assumption, rather than a measured regional loss.

Rent duration: operating margin falls toward 46.0% from FY2036, using the FY2022–FY2026 historical reference, rather than perpetuating today's scarcity. Earlier normalization in FY2032 costs $8.81 per share; the separate 58% long-run-margin hypothesis adds $39.17. Neither the completed AI cycle nor a permanent margin is established by that history.

Reinvestment and cash: FY2028 operating cash of $351.4bn leaves approximately $319.0bn after physical investment and the economic compensation adjustment; financing-neutral firm cash flow is $317.2bn. Receivables and inventory consume cash as sales rise. H1 FY2027 investment gains of $23.7bn do not recur as an operating engine; the FY2028 tax assumption is 18%.

What the market already expects: the compiled revenue estimates are $411.7bn for FY2027 and $692.2bn for FY2028, close to this forecast. The snapshot does not establish a broad overlooked growth story; its earnings-per-share accounting basis was unverified. Greater exposure needs a better accessible-demand or rent path, judged against Section 10's boundaries.

10. Growth is visible; price leaves little room for error

The current financial calculation is dated 6 November 2026 and uses that session's $186.64 close. Expected intrinsic value is $206.07; carrying it at the 12.60% equity hurdle and deducting the assumed $1 dividend gives $231.04 for 6 November 2027. Total return is 24.32%. Financial BUY requires more than 22.60%, SELL less than 2.60%, with equality in HOLD. The arithmetic signal is BUY; the business-quality and realization requirements remain open, and no active recommendation is assigned.

Required returns and valuation

The cash-flow valuation runs through FY2041 and then uses 3% nominal terminal growth. At 6 November inputs, the base operating discount rate is 12.5172%; non-operating assets and liabilities enter once, and each outcome is divided by its fresh forecast diluted claims before the 55/30/15 base/bear/bull weighting. The equity hurdle is 12.60%, derived from the 4.86% Treasury observation, retained 4.50% premium and

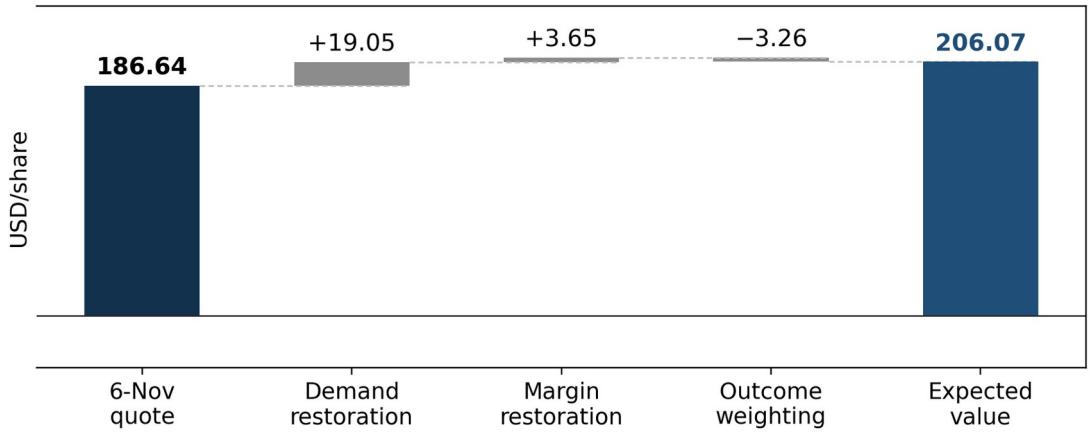
source beta 1.72. Cash-flow fractions, discount periods and the cash/asset/share bridge use the 6 November valuation date.

The decision quantity is the uniform adjustment to annual Hyperscale and Industrial demand growth in FY2029–FY2031. At the unchanged central 0pp, return exceeds the financial BUY threshold by 1.72pp. Full inverse runs put the BUY boundary at -0.73pp and the SELL boundary at -9.82pp, holding other inputs and outcome weights fixed. Results and outlook on 17 November can test delivery and demand; a quarterly beat alone does not prove the later path.

The quote is reproduced when the same later-year demand adjustment is -5.11pp, with other assumptions fixed. This is a conditional economic requirement, not a measurement of investors' beliefs. The $186.64 entry price leaves room for a weaker operating path than the central forecast; it supplies no evidence that customers will keep buying at exceptional margins.

Figure 6 starts from a price-reproducing alternative under the same 6 November financial inputs. With later demand growth 5pp below the central path, a solved -0.88pp operating-margin shift reproduces $186.64 in the base case. Restoring central demand adds $19.05; restoring central margin changes value by $+3.65; applying the existing outcome weights changes it by $-3.26, ending at $206.07. The construction keeps each economic step visible without claiming that it recovers the market's actual forecast.

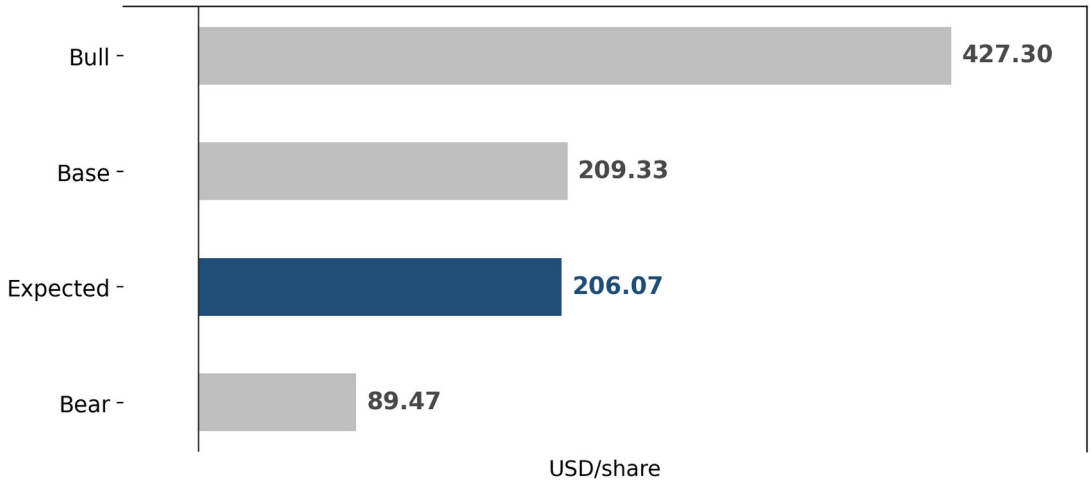
Figure 6: Demand, margin and outcome weighting take the 6 November quote of $186.64 to $206.07 expected value. The financial BUY/SELL growth boundaries are -0.73pp / -9.82pp.



Source: Separate complete 6 November 2026 formula-graph runs; quote-reproducing operating alternative, central path and unchanged outcome weights.

The central value conceals a wide outcome range. The weaker path combines poorer delivery, demand and retention with earlier margin normalization; the stronger path combines expansion with longer-lasting rent. Figure 7 and Table 6 show those outcomes:

Figure 7: The retained 30/55/15 bear/base/bull weights produce $206.07 expected value. The current financial hurdle is 12.60%; the return band is ±10pp.



Source: Separate complete 6 November model outcomes, each divided by its own forecast diluted shares before weighting. Probabilities remain an explicit analyst judgement.

Table 6: The 30/55/15 weights produce $206.07 expected value; the $209.33 base case is a different object.

Case	Value today $	12m target $	12m total return %	Weight %
Bear	89.47	99.74	-46.02	30
Base	209.33	234.71	26.29	55
Bull	427.30	480.14	157.79	15
Expected	206.07	231.04	24.32	—

Source: Current 6 November full outcome calculations. Probabilities are analyst judgements; expected value is a weighted aggregate rather than a fourth scenario. The $1 ordinary dividend is assumed. The bear case is not a liquidation floor.

The multiple cross-check tests a different assumption. At the current $186.64 price, estimated normalized FY2027 earnings of $9.19 imply about 20.30 times earnings. Applying analyst conventions of 15, 20 and 25 times gives values of $137.90, $183.87 and $229.84. This span does not independently reproduce or validate the discounted-cash-flow value. It shows how much the answer depends on the price paid for a given earnings stream; the multiples are chosen illustrations, not an observed comparable-company range.

The three-year illustration starts at the $186.64 close on 6 November and ends on 6 November 2029. Fiscal FY2027 and FY2030 earnings are retained as the comparable earnings proxies, rather than an exact quarterly exit forecast. Normalized earnings of $9.19 grow to $22.11; a fifteen-times exit multiple gives $331.62 and 21.57% annual return including $1 dividends each year. Ten and twenty times instead give 6.31% and 33.72%. At fifteen times, the exact annual contributions are 0.45pp paid cash, 32.34pp per-share growth and -11.23pp rerating. The separate FY2028 available owner-cash yield is 7.07%; it is not added again as paid income. This assumed exit regime does not activate a recommendation.

Value reaches the owner through useful systems delivered, customers paying and earnings remaining after reinvestment. The financial BUY signal at the 6 November entry price does not itself establish business quality or a realization mechanism. Chinese deployment evidence has not been calibrated into a measured regional forecast. Displaced paid workloads, lasting price concessions or a shorter high-margin period would require revised inputs and a full valuation.

11. The next evidence must reach deployment and cash

NVIDIA's next dated test is the **17 November 2026** Q3 release, scheduled at approximately **13:20 Pacific**, followed by the **14:00** call; the quarter ended **25 October**. Orders and deployment plans are early signals. Independent buyer installations and useful workload economics confirm the mechanism. Reported revenue, receivables, operating cash and distributions establish the later result. A share-price reaction is not another operating observation.

Deliveries and an updated demand anchor can justify a stronger later path; order losses, poor utilization, delayed installations or slower collection can justify a weaker one. Current financial boundaries are -0.73pp for BUY and -9.82pp for SELL on the named later-year demand adjustment. These are path-wide inverse conditions, not single-quarter company targets. An absence of new long-range guidance does not confirm the existing forecast.

Huawei has announced 30 November 2026 for AICS availability outside China; Alibaba has placed V900 commercial release in Q1 2027. For these competitive tests, an announcement becomes evidence of availability only when a customer can obtain the service or system. Repeat paid deployment at comparable quality, response time and full cost would be stronger confirmation; a roadmap date or token count alone would not establish lost NVIDIA orders. These observations test the competitive mechanism in Section 5, rather than an assigned regional haircut.

FY2027 ends **31 January 2027**; an unannounced results-publication date is not assumed. Subsequent balance and cash statements must distinguish delayed collection from recurring economic leakage. We are wrong to dismiss a shortfall as timing if receivables, inventory and noncash adjustments show persistent leakage. The priced transfer alternatives and their limits remain in the appendix.

12. What can turn the operating success into capital loss

Permanent loss begins when NVIDIA's platform no longer earns enough premium to cover the cost of sustaining it. Customers can migrate workloads, competitors can make migration economical, and buyers can gain negotiating power as capacity expands. Retained orders and pricing then fall while commitments remain. The weaker joint outcome is worth $89.47 per share with -46.0% twelve-month return at the 6 November entry price; it is an economic scenario rather than a floor. Sections 7–9 separate retention, capture, reinvestment and rent-duration channels.

Technology leadership must survive the customer's workload. The matched Rubin evidence is a preview with modelled costs and a token mix dominated by input and cache. Its roughly 1.90-times tokens-per-dollar advantage requires approximately 53% relative useful utilization under the stated conditions; a tighter service requirement can reverse the comparison. Actual purchase prices, deployed utilization and NVIDIA's capture of the saving remain unestablished. A benchmark result therefore cannot justify a universal demand elasticity or replacement cycle.

Supply, deployment and export access can fail together. Packaging, memory, power or geopolitical disruption can delay systems and leave contractual exposure, even after a buyer commits capital. NVIDIA-specific packaging and memory allocations are not measured. A small cash-timing experiment does not price a prolonged interruption: deliveries, margin, customer solvency and obligations could all change together. Product restrictions can also reduce accessible demand or cause inventory and contract losses. Direct-customer headquarters do not identify ultimate end demand. China's alternative chain can also serve customers through a cloud service outside China. The excluded near-term compute outlook in Section 6 therefore does not bound the competitive exposure described in Section 5.

Investment assets, guarantees and weak collection supply further routes to permanent loss. Newly purchased equity may subsequently be impaired; conditional guarantees may crystallize if projects fail; accounting profit may arrive before cash. The technical appendix retains the separate priced asset haircuts, guarantee loss, receivables delay, noncurrent-asset classifications and temporary/permanent cash-transfer alternatives. They are distinct hypotheses, rather than amounts to sum into a new outcome or proof that every exposure has a current market price.

Quality still requires evidence across a complete cycle, sufficient maintenance, an advantage in the worst year and money reaching continuing owners. The limits set out in Section 8 remain independent of the calculated target. A financial BUY threshold alone would not prove those conditions.

The stated unrestricted global institutional investor can execute through the listed shares, subject to its legal access, tax and trading arrangements. Values are in US dollars before investor-specific tax and dealing costs. Declared and paid dividends are cash to the owner; retained investments and future buybacks reach that owner differently.

Finally, scenario probabilities and the equity premium are analyst judgements; the price and beta use a preserved market-data relay. The external estimate snapshot is incomplete, its earnings basis unverified, and terminal growth is a convention rather than a measured AI life cycle. The resulting uncertainty needs specific new operating evidence and priced alternatives. It does not authorize an arbitrary wider value range or an unchanged investment conclusion after new evidence.

13. Sources

Valuation and the market reference are dated 6 November 2026; the reference is the completed regular session at 16:00 EST. The target horizon is 6 November 2027. Operating-source dates and their disclosed periods retain their own meaning; the information set for the strategic China discussion remains identified below. Fiscal years are NVIDIA fiscal years unless a buyer period is named. Reported, estimated and conditional quantities remain distinct.

1. NVIDIA FY2026 Form 10-K, year ended 25 January 2026: consolidated income and cash flows, segment reconciliation, market platforms, property and equipment, commitments and equity. Historical FY2022–FY2025 comparisons also use FY2022 Form 10-K, FY2023 Form 10-K, FY2024 Form 10-K and FY2025 Form 10-K. Figures use consolidated rather than unadjusted segment profit unless stated.

2. NVIDIA Q1 FY2027 Form 10-Q, quarter ended 26 April 2026, and Q2 FY2027 Form 10-Q, quarter and half-year ended 26 July 2026: current operating statements, balance, cash flows, reportable segments, recast platforms, direct-customer concentration, assets and obligations. Current platform comparisons follow the Q2 classification.

3. NVIDIA Q2 FY2027 CFO commentary and official 26 August 2026 earnings-call transcript: Q3 outlook, FY2027 tax expectations and preliminary FY2028 growth/margin comments. Guidance is distinguished from reported results and the later analyst forecast.

4. TSMC Q3 2026 call transcript, particularly the capital-expenditure, accelerator and advanced-packaging discussion on pages 4, 7, 12 and 17. It does not establish NVIDIA-specific absolute packaging or memory allocation.

5. Buyer investment and composition: Microsoft Q1 FY2027 results and call; Alphabet Q3 2026 release and 28 October call; Amazon September 2026 Form 10-Q; Meta Q3 2026 release. The cash-PPE comparison uses the stated common accounting perimeter.

6. External custom-chip competition: Alphabet's call above and Amazon's original discussion of custom-chip demand and outside sales. These disclosures establish a competitive route, not a measured percentage of displaced NVIDIA orders.

7. Technology comparisons: MLCommons MLPerf Inference v5.0 and NVIDIA's original per-system/per-GPU discussion, 2 April 2025. InferenceX original AgentX results, 19 October 2026: only the matched DeepSeek-V4-Pro 100-tokens/s/user comparison supports the current numerical claim. TCO is modelled and preview/software/token-mix limits remain material.

8. Figure 2: NVIDIA Technical Blog, “Figure 1. NVIDIA Blackwell Ultra GPU chip explained”, 22 August 2025, updated 24 September 2025. This is a borrowed manufacturer architecture illustration, separate from the six quantitative research charts.

9. Worst-year comparison: AMD calendar-2022 results, consolidated operations and acquisition-amortization reconciliation. Product mix, acquisitions and different fiscal end dates limit the comparison.

10. Market inputs: US Treasury daily par yield curve, 6 November 2026 ten-year CMT 4.86%, observed near 15:30 New York; the Nasdaq explicit regular-session Closed reference of $186.64, crosschecked to daily close and excluding after-hours quotes. The beta 1.72 is the measured 260-week Nasdaq-100 estimate, not a newly resampled 6 November beta. The 4.50% equity premium is the common analyst convention, not a market observation.

11. External estimate comparison: Yahoo Finance NVIDIA analysis, observed 2 November 2026 through search extraction. The revenue sample sizes are retained; full-page access was rate-limited and the selected EPS basis was not established. These compiled estimates do not prove the absence of a view from all market research. Analyst exit multiples and 30/55/15 outcome weights are explicit conventions of this analysis, not values taken from that snapshot.

12. Calendar: NVIDIA's official Q3 FY2027 results-call notice, published 28 October 2026, announces the 17 November results and call. The FY2027 end date follows the disclosed 53-week fiscal calendar; it is not an assumed earnings-release date.

Chinese models and the competing computing chain

The strategic analysis uses an information cutoff of 2 November 2026. The following primary materials support the discussion; company-described tests and commercial claims are identified as such in Section 5.

- DeepSeek V4.1 technical report, September 2026, pp. 1 and 4–5; version-pinned report linked in Section 3. DeepInfer release, 20 October, pre-cutoff revision 25 October; both NVIDIA and Ascend support, with performance plots on NVIDIA hardware.
- Qwen3.8-27B version-pinned model card and Apache 2.0 licence. Qwen Max uses different commercial terms; no claim is made that every Qwen model is unrestricted.
- Alibaba June-quarter 2026 results, 20 August, pp. 3–4; its international cloud and V900 update, 25 September. Cloud revenue and region counts do not measure NVIDIA substitution.
- Huawei and SiliconFlow, CloudMatrix study v2, 18 June 2025, §§3.3 and 5.1–5.3. Huawei AICS announcement, 18 September 2026; the 30 November overseas date is a plan at the cut-off.
- SMIC interim report, 3 September 2026, pp. 7 and 13–14; JCET interim report, 21 August 2026, p. 12. Aggregate fabrication and packaging disclosures do not establish AI-chip yield, memory availability or an attributed Ascend supply chain.

Valuation bridge at 6 November 2026

The valuation uses the 6 November 2026 regular-session close and a twelve-month horizon ending 6 November 2027. The arithmetic financial signal is separate from an active investment recommendation.

Paired calculation: NVIDIA (NVDA) — модель.xlsx

Calculation identity: complete 6 November formula graph SHA-256 7b6341aff73f330d6bd69e23435c976ab028ef11e37cfc01c7b1df92df4f4afa. The paired current workbook and its exact file fingerprint are listed in the delivery manifest.

The operating assumptions and source classifications are retained. All following bridges, outcome values and sensitivities use full dependency replays at 6 November financial inputs. The China discussion remains illustrative; it supplies no measured regional forecast.

BUY is the current probability-weighted financial signal: 24.32% total return against a 12.60% hurdle and a ten-percentage-point band. All four business-quality tests remain unestablished. Tables 7 and 8 show the active base case; Table 9 reconciles it to the expected value on the first page.

Table 7: Base explicit cash flows contribute $3,393.7bn of present value. USD bn.

Fiscal stage	Remaining FCFF	Present value
Remaining Q4 FY2027	49.4	48.1
FY2028	317.2	274.2
FY2029–FY2031	1,388.1	944.1
FY2032–FY2036	2,994.4	1,284.4
FY2037–FY2041	3,546.0	842.9

Source: Paired workbook, Valuation cash-flow and present-value rows. Stage totals sum the individual annual flows; present values use each period's own date. SBC remains an expense.

The base discounts firm cash flow at 12.5172%. Final FY2041 cash flow grows by 3.0% to $776.4bn and is capitalised at WACC less growth. Terminal present value is 30.96% of enterprise value. These are base-case amounts rather than probability-weighted enterprise values.

Table 8: The current base bridge gives $209.33 per share before outcome weighting. USD bn except shares and per-share values.

Component	Value
PV explicit cash flows	3,393.7
PV terminal cash flows	1,522.1
Enterprise value	4,915.9
Cash and debt securities	80.1
Debt unpaid consideration and tax reserve	-33.8
Existing public / private asset value	47.7 / 38.4
New funded non-operating assets	10.1
Base equity value	5,058.3
Estimated diluted shares bn	24.1639
Base V0 → TP12; D12	$209.33 → $234.71; $1.00

Source: Paired workbook, Valuation. No future allocation loss or guarantee loss is assumed centrally; separately priced loss stresses are supplied in the accompanying calculation package, without treating unmeasured risk as zero.

Table 9: The retained 30/55/15 weights give $206.07 expected value and $231.04 at the twelve-month horizon.

Outcome	Weight	V0	TP12	TSR
bear	30%	$89.47	$99.74	-46.02%
base	55%	$209.33	$234.71	26.29%
bull	15%	$427.30	$480.14	157.79%
Expected	Weighted	$206.07	$231.04	24.32%

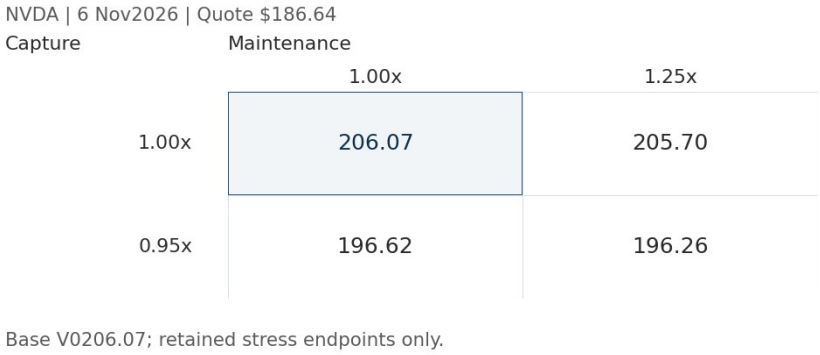
Source: Complete 6 November case calculations, with fresh scenario-specific claims. TP12 + $1.00 = V0 × 1.126; TSR uses the $186.64 close of 6 November. Case probabilities are the existing judgement.

Rounding: amounts in USD billions are displayed to one decimal. Percentages, ratios and per-share values are calculated before rounding; totals may differ from the sum of displayed rows.

Sensitivity on the 6 November valuation

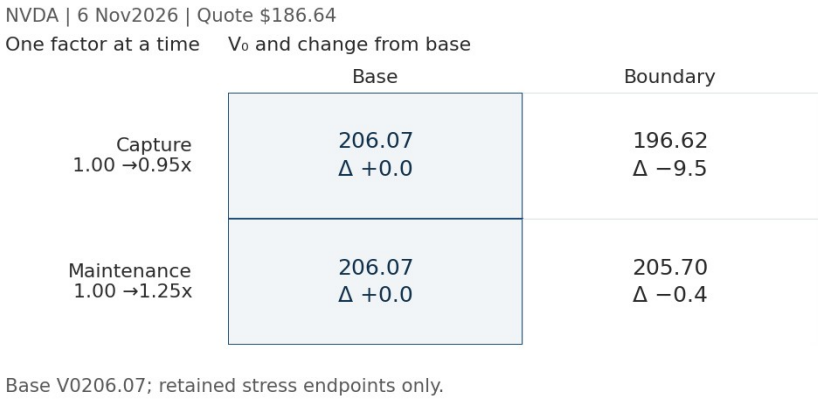
Expected value retains 55% base, 30% bear and 15% bull weights. Each endpoint changes the named operating input in a complete current-date run; expected-value effects differ from active base-case effects elsewhere in the note.

Figure 8: The joint retained stress reduces V_0 from $206.07 to $196.26.



Source: Complete 6 November NVIDIA formula graph, quote $186.64, Rf 4.86%, ERP 4.50%, target horizon 6 November 2027. Clean separate dependency replays at retained operating boundaries.

Figure 9: Separate operating boundaries change V_0 by $-9.46 / $-0.37.



Source: Complete 6 November NVIDIA formula graph, quote $186.64, Rf 4.86%, ERP 4.50%, target horizon 6 November 2027. Clean separate dependency replays at retained operating boundaries.

Probabilities remain base 55%, bear 30%, bull 15%. The two retained endpoints per factor are stated assumptions; no unsupported interior or opposite boundary is asserted. Current V0 and all claims are replayed for each cell. The financial signal is not an active recommendation; native Excel and financial-release acceptance are reported separately in the delivery evidence.

Cross checks sources and material limitations

The useful cross-check is the cash and growth requirement behind the price. With outcome weights and other assumptions fixed, the FY2029–FY2031 demand adjustment is 0pp; the financial BUY boundary is -0.73pp and SELL -9.82pp. This is an inverse of the same economics, not an independent target. The 58% long-run margin alternative raises base value from $209.33 to $248.51; the $39.17 difference shows why segment mix alone cannot justify additional rent.

The preserved external estimate snapshot puts FY2027/FY2028 revenue at $411.7bn/$692.2bn against approximately $410.7bn/$698.1bn in the forecast. Near-term revenue is therefore not the main disagreement. The snapshot is a secondary compilation observed 2 November, with an unverified EPS basis; it is not used to assign an independent valuation weight.

Table 10: Material inputs combine issuer disclosures with explicit dated estimates.

Input and period	How obtained and source	Material boundary
Q2 and H1 FY2027 revenue and segment profit	Disclosed; NVIDIA Q2 FY2027 10-Q, quarter ended 26 July 2026	Segment profit and platform revenue are distinct
FY2027 outlook and FY2028 growth anchor	Disclosed management expectations; 26 August CFO commentary and call	Guidance is not delivered revenue; later growth and rent fade are estimated
Rf 4.86%; beta 1.72; ERP 4.50%	Rf: disclosed Treasury 6 November; beta: 260-week estimate; ERP: common convention	Ke 12.60% calculated; beta is not a new 6-Nov resample; price is Nasdaq Closed reference crosschecked to daily close
Investment assets and scenario probabilities	Disclosed July marks with estimated haircuts; existing 30/55/15 outcome judgement	Marks are not current exit prices; no new probabilities inferred

Source: Original issuer disclosures and operating assumptions, with 6 November price, financing and valuation-date inputs. Exact formulas and input provenance accompany the paired calculation.

The primary Q2 FY2027 disclosure confirms $89.0bn of Data Center market-platform revenue ($89,023mn), versus $88.3bn of Compute & Networking reportable-segment revenue ($88,299mn). The $724mn difference reflects different classifications. Data Center plus Edge Computing, and Compute & Networking plus Graphics, each reconcile separately to the same $96.2bn consolidated revenue. The two measures must not be substituted or added together.

NVIDIA Q2 FY2027 Form 10-Q — reportable segments and revenue by market platform

Table 11: The next evidence must resolve deployment cash and persistent rent.

Unresolved item	Exposure or measured alternative	Next observation and owner
Useful workload economics and capture	One-time capture −5% costs $9.60/base share; China-specific displacement is not measured	Analyst: matched deployed workloads and realised pricing during FY2027–FY2028
Margin durability and cash conversion	58% alternative adds $39.17/base share; FY2026 CFO residual $4.1bn is not a probability band	Analyst: 17 November results and later cash receipts; compare like-for-like margins
Assets market data and external estimates	July asset marks and unverified exchange tape/EPS basis; aggregate valuation effect unmeasured	Analyst: next investment disclosure and independent price/estimate records before a future revaluation
Business quality and recommendation status	Four-test quality evidence remains unestablished; no active recommendation assigned	Research owner: complete cycle maintenance advantage and minority distribution evidence before new issuance

Source: Existing separate full-model stresses and dated observations in Sections 7–13. The displayed effects are not additive and do not price every missing piece of evidence.

Reporting segment results at the retained disclosure

Two full financial years and the latest disclosed comparable period are shown on the original source basis. Revenue growth compares Q2 FY2027 with Q2 FY2026; operating profit and margin refer to Q2 FY2027. USD bn, except percentages.

Table 12: Reported businesses reconcile to consolidated revenue and operating profit.

Reporting segment	FY2025 revenue	FY2026 revenue	Q2 FY2027 revenue	Revenue YoY	Q2 FY2027 op. profit	Op. margin
Compute & Networking	116.2	193.5	88.3	113.6%	62.7	71.0%
Graphics	14.3	22.5	7.9	46.4%	3.9	49.2%
Corporate reconciliation	—	—	—	—	−2.9	—
Consolidated NVIDIA	**130.5**	**215.9**	**96.2**	**105.9%**	**63.7**	**66.2%**

Reporting segments differ from market platforms such as Data Center. Corporate costs reconcile segment profit to consolidated operating income; platform revenues are shown separately in the existing detail.

“—” means not applicable or not separately reported; it is not a zero estimate. Revenue growth and operating margins are calculated from the disclosed amounts before rounding. Operating profit is the issuer’s segment measure, reconciled through the central row where applicable.

Source: Issuer Form 10-K (FY2025 and FY2026) and Form 10-Q (Q2 FY2027); reporting-segment notes. Annual and interim periods retain their distinct horizons.

NVIDIA annual report — FY2026

NVIDIA interim report — Q2 FY2027