Ed Faritov, CFA

Market	US; AI
Company	Anthropic PBC
Ticker	ANT, proposed
Status	Pre-IPO
Price and target	No public quote or target
Rating	Not assigned
Series H	$965.0bn; 28 May 2026

Anthropic and the Economics of Paid Outcomes

Report dated 7 November 2026 | Information cut-off 5 November 2026

Anthropic sells completed work through models and applications. Its central risk is the mismatch between a fast repricing of useful intelligence and years of compute payments. Verification can defend a premium, but spending share does not prove pricing power. Gross margin excludes much of frontier upkeep. We assign no investment rating or public price target before final IPO terms. S-1/A.

The premium rests on verifying the outcome. For the specified 100,000 input and 10,000 output tokens, Sonnet costs $0.30 and DeepSeek Flash $0.021 at off-peak rates. Ten seconds of review at $100 an hour absorbs the difference. A cheaper model that needs no extra oversight erodes the premium quickly, even when infrastructure payments continue. This is a test, not measured productivity. Anthropic, DeepSeek.

Efficiency helps the owner only after customers, usage and contracts take their share. A compatible software improvement can use installed chips without a new architecture purchase. More paid work can fill the freed capacity, while price cuts pass savings to customers. The $69.8bn of near-term infrastructure obligations does not fall automatically with compute per task.

Operating evidence is stronger than the net loss suggests. Audited annual 2023–2025 statements and unaudited nine-month 2026 statements are public in the 2 November S-1/A. Nine-month gross margin was 85.5%, research was 57.9% of revenue and the operating loss was $1.3bn. Cash remaining after capital and lease payments was $3.6bn, before normalization for compensation, working capital and capacity timing. Financial statements.

The historical $965.0bn Series H mark demands both scale and a durable residual margin. At the retained 15% return, 20% ownership loss and 25 times cash flow, a 25% cash margin requires $388.2bn of revenue in five years. With a 20% margin and 30% ownership loss together, the requirement becomes $554.6bn. These are conditional entry diagnostics, not a current private quote or public target. Series H.

1. How the Market Creates Demand
2. Products and the Economics of Outcomes
3. Competitors and Concentration
4. Strategy: Capacity, Distribution and Constraints
5. Financials and Owner Earnings
6. Funding Rounds and Reverse Valuation
7. Scenarios, Observations and Risks
8. Sources

1. How the Market Creates Demand

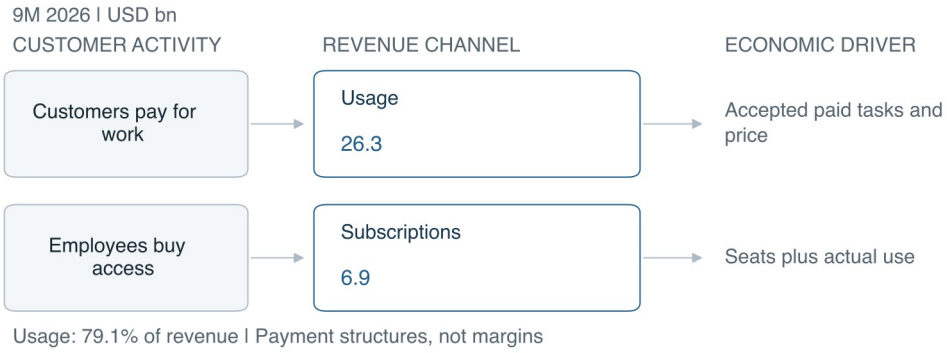
The AI market is growing around work tasks, but several participants share the proceeds. Developers buy access to a model through an application programming interface, or API; employees buy Claude or Claude Code; a business may buy a completed legal document or a resolved support ticket. Infrastructure supplies compute, while the application connects the model to data, permissions and rules for accepting the outcome. Growth in one layer does not pass all of its economic value to the others.

Menlo estimated the annualized pace of enterprise spending on deployed AI at $176.8bn. This is an estimate from a survey of almost 400 US executives covering more than $4.0bn of spending, not industry revenue for a calendar year or a ceiling on the market. Its estimate for coding applications is $34.0bn, of which $24.0bn comes from the labs' own applications. Menlo is an investor in Anthropic; its findings are useful as a measurement of a specific sample. Menlo, 3 November 2026.

The economic source of demand is the cost of human work. However, the wage bill is not the labs' addressable revenue: some tasks cannot be automated, some savings stay with the customer, and some go to the application and integrator. If the price of an outcome falls, more paid outcomes are needed to sustain revenue. Anthropic's claimed 40% reduction in the cost of a typical Opus task would require 66.7% more such tasks merely to maintain revenue, all else equal. This is conditional arithmetic; actual task pricing and volume for a comparable customer cohort are not disclosed. Opus 5.5.

Usage remains Anthropic's principal source of revenue. In the first nine months of 2026, usage generated $26.3bn and subscriptions $6.9bn; usage accounted for 79.1%. Chart 1 shows where the customer decides to pay. Cloud marketplace sales are a channel: their 44% share cannot be added to usage and subscriptions. For the owner, this means expanding use matters more than the number of seats sold alone. S-1/A, Note 2.

Chart 1: 79.1% of revenue is paid for by usage

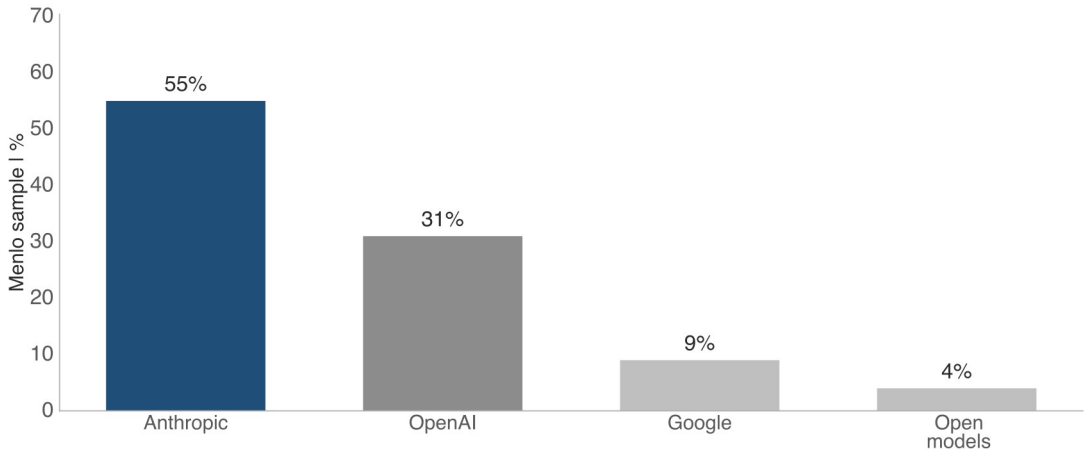


Value sensitivity: Per-share value awaits final IPO capitalization

Source: S-1/A, revenue for the first nine months of 2026. Sales channel and payment structure overlap.

Anthropic's share of API spending reaches 55% in Menlo's sample (Chart 2). This measures a supplier's dollar share in the survey; Section 3 compares it with the customer's ability to choose several models.

Chart 2: Anthropic captures 55% of API spending in Menlo's sample



Source: Menlo, September 2026 sample. Published shares sum to 99%; no separate interpretation is assigned to the remainder.

2. Products and the Economics of Outcomes

Claude sells a choice between the cost of an attempt and the cost of an error. Opus is designed for complex work, Sonnet offers a less expensive balance of quality and speed, and Haiku serves high-volume processing. Claude Code adds work with repositories, tools and tests. Subscriptions sell an interface and access, while Enterprise combines seat charges with usage. Successful application growth can therefore come with variable revenue that depends on activity. Anthropic products, Claude pricing.

The product map represents four distinct customer buying decisions. A customer may simultaneously buy a model for its own application and Claude seats for employees. These expenditures solve different problems; the model's technical quality does not replace the economics of the finished product.

Table 1 distinguishes the products by purchase, source of value and boundary of responsibility.

Table 1: Where Anthropic meets the buyer and what it is paid for

Product	Buyer and payment	Source of value	Constraint
Models and Claude Platform	Developer; API and tool usage	Intelligence and components for the customer's own application	The customer builds the workflow and accepts the outcome
Claude Code	Developer and team; plan and usage	Work on code, tests and repository changes	Access permissions and review of changes are required
Claude and Cowork	Employee and manager; subscription and usage	Work with files, applications and multistep tasks	Errors move from text into actions
Enterprise	Corporate customer; seats and usage	Access, security and deployment management	Procurement approval does not yet establish usefulness

Sources: Claude Platform, Claude Code, Cowork, Enterprise, pricing. This is a map of offerings, not disclosed financial segments.

Platform sells developers building blocks; Code sells a connected execution workflow. With Platform, customers receive direct access to the model or a managed agent environment: they can retain their own interface, work queue and relationship with the end buyer. Code moves some of this assembly into Anthropic's product: planning,

editing and testing take place alongside the repository. The economic difference is who pays for integration and who controls the next choice of model. Platform, Code.

Cowork extends the audience beyond developers. Its purpose is to complete a task in selected files and tools while allowing users to direct and review the work. The company combines Cowork with conversational Claude: the transition from asking a question to execution becomes part of one application. This opens up finance, research and administrative workflows, while also raising the cost of an incorrect action. Reviewing finished text and authorizing an agent to change business systems are separate customer decisions. Official Cowork page.

Enterprise turns a useful personal tool into an authorized corporate workflow. Security, administration and connections to internal systems reduce friction in expansion: an individual enthusiast is insufficient if procurement cannot grant employees access. Defensibility can arise from the combination of permissions, tested procedures and team training. The data nevertheless remain the customer's asset; Anthropic states that corporate data are not used for training by default. Established access therefore cannot automatically be treated as a proprietary, unique training dataset. Enterprise.

Commercial demand has already moved beyond isolated experiments. As of 30 September, Anthropic had more than 11,900 customers with at least $100,000 of recognized revenue over the preceding twelve months, versus more than 950 a year earlier; more than 1,900 customers generated at least $1mn. The company aggregates accounts belonging to one organization: this measures revenue per customer, not user numbers. This scale establishes a willingness to pay. The next proof is retaining and expanding these relationships after initial deployment. S-1/A, customer metrics.

Table 2: Consumption revenue reached $26.3bn in nine months; profit by revenue channel is not disclosed.

Revenue and operating result in USD bn. Current period: 9M 2026; revenue YoY compares with 9M 2025.

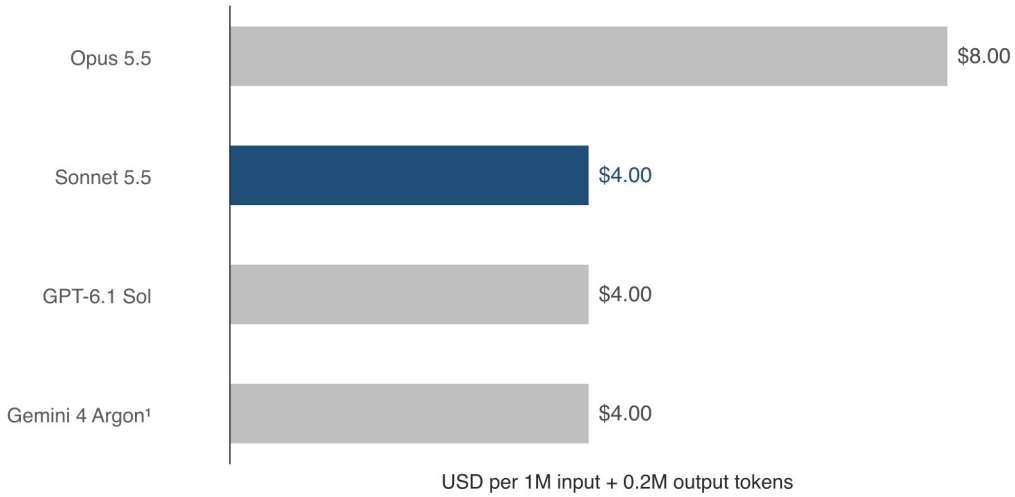
Reporting segment / reconciliation	Revenue FY2024	Revenue FY2025	Revenue 9M 2026	Revenue YoY	Operating result[1]	Operating margin[1]
Consumption-based revenue	0.3	3.8	26.3	1137.9%	n.d.	n.d.
Subscriptions	0.1	0.8	6.9	1551.4%	n.d.	n.d.
Consolidated Anthropic	**0.4**	**4.6**	**33.2**	**1206.2%**	**−1.3**	**−3.8%**

[1] Operating result and margin refer only to the current period; margin is operating result divided by revenue. One reportable segment. Consumption and subscriptions are revenue channels, not separate operating segments; their profit is not disclosed. The total operating result is consolidated operating loss, not the CODM measure of consolidated net loss. n.d. means not disclosed. The total operating result is a consolidated measure, not the CODM segment measure of consolidated net loss.

Sources: Anthropic Form S-1/A, consolidated financial statements and Note 14. Official filing 1

Sonnet 5.5, GPT-6.1 Sol and Gemini 4 Argon have the same standard price as of the research date: $2 per million input and $10 per million output tokens. A basket of one million input and 200,000 output tokens costs $4; for Opus it costs $8. Argon's rate is introductory through 30 November: the published rate doubles from 1 December. Chart 3 compares only this basket; caching, long context, additional tools and different token counts change the actual bill. Sonnet, OpenAI, Google.

Chart 3: Sonnet, GPT and Gemini each cost $4 for the same pricing basket



Sources: Anthropic, OpenAI, Google. As of 5 November 2026; Argon uses the introductory price through 30 November 2026.

Our first conclusion: savings on verification defend Claude's premium, rather than the nominal price of intelligence. For 100,000 input and 10,000 output tokens without caching, using DeepSeek's off-peak rate, the difference between Sonnet and DeepSeek Flash is $0.279. Ten seconds of labor at $100 an hour costs the same. We are wrong for this task if the cheaper model maintains the share of accepted outcomes and adds less than 10 seconds of total review and correction. Token counts, the wage rate and an identical success rate are specified for this test; this calculation does not establish a measured advantage for Claude. DeepSeek pricing.

A reduction in inference work per accepted task is economically valuable even on the same installed chips. It can reduce short-term top-up purchases or support more paid tasks on committed capacity. It does not prove that a new architecture has been bought, delivered or made productive. If growth in useful paid tasks more than offsets the reduction in compute per task, total compute use can still rise. Faster inference is therefore neither a measured decline in total compute demand nor a dollar-for-dollar cash saving.

The entire path to an accepted outcome must be tested: attempts, tools, rework and human effort. In a 2025 experiment, METR found that experienced developers took 19% longer; that finding concerns older models and specific repositories, not Claude 5.5. It illustrates the cost of measurement error: a subjective sense of acceleration can diverge from the facts. For shareholders, customer retention after such a complete assessment is the evidence of paying demand. METR primary research.

3. Competitors and Concentration

Anthropic leads on one dollar measure of the market while facing strong rivals. OpenAI competes through models and workplace applications; Google combines models with its own distribution and infrastructure; specialized applications control the customer's work queue. DeepSeek and downloadable models offer an alternative cost of execution. These rivals affect different parts of the customer's bill, so neither price lists nor benchmarks yield a single ranking of winners.

Table 3 compares competitors' established strengths with Anthropic's potential response.

Table 3: Competitors challenge different parts of Anthropic's offering

Competitor	Established strength	Pressure on Anthropic	Strategic response
OpenAI	Models, ChatGPT and Codex	Workplace interface and API budget	Connect Claude's quality with everyday work
Google	Gemini, applications and cloud	Distribution and existing services	Multiple clouds and proprietary applications
Microsoft	Workplace applications and enterprise context	Interface, data and procurement	Integration and a direct channel
Open-weight models and DeepSeek	Hosting choice or low API prices	Price of sufficiently good execution	Total cost of an accepted outcome
Cursor and Cognition	Specialized development tools	Work queue and relationship with the team	Code and the appeal of Platform

Offering sources: OpenAI, Google, Microsoft, DeepSeek, Meta, Cursor, Cognition. The final column is our interpretation of Anthropic's strategy; the table does not rank model quality.

An established workflow offers better protection than leadership on a single benchmark. Our strategic view is that a durable Anthropic advantage must combine model quality, access to the task and accumulated knowledge of how the customer accepts an outcome. Proprietary evaluations, user corrections and an understanding of permitted actions can improve the application and its assessment. But this feedback does not automatically create a monopoly: customers may own the evaluations and provide the same context to a competitor. The strongest rival is whoever already sits between the employee, the data and the button that accepts the work.

Proprietary applications create a conflict with API customers. Code and Cowork allow Anthropic to get paid closer to the outcome and depend less on someone else's interface. They also compete with developers building their own applications on Claude. If the lab implements the same functions itself, a Platform customer has an incentive to keep a backup model. Anthropic's response is to keep the platform's quality, availability and tools attractive even where it does not own the application. Otherwise, success in the direct channel could weaken distribution of the model through partners. This is a strategic risk of its dual role, not established customer attrition.

In Menlo's sample, Anthropic receives 55% of API spending, OpenAI 31%, Google 9% and open models 4%. Published shares sum to 99%; a separate category for the remainder is not established. The three named closed-model providers alone contribute 4,067 points to the concentration index, calculated as the sum of squared shares. No comparison with other markets on a consistent basis is available. Chart 2 represents API dollars, not users, tokens or the companies' total revenue; the survey predates Argon's broad availability on 13 October. Menlo, Google launch.

Our second conclusion: concentrated spending does not yet justify assuming durable pricing power. In a16z's survey, 78% of respondents used OpenAI in production systems and 44% used Anthropic; 81% tested or used at least three model families. Ramp, by contrast, recorded Anthropic payments from 43.8% of its customer panel in September, versus 40.6% for OpenAI. The differences show that results depend on the population and unit of measurement. None of these shares measures the response to a price increase. a16z, Ramp.

Our caution on pricing power is wrong if the actual price rises 10% while Anthropic's share of comparable paid tasks holds for two quarters without offsetting discounts. Additional defensibility may reside in configurations, security checks and customer data. But DeepSeek already describes how to connect its own model to Claude Code: attachment to the tool is technically separable from model spending. For shareholders, retained margin after that substitution is what matters. DeepSeek instructions.

Chinese model quality and Chinese chip capability require separate evidence. Domestic hardware sufficient for local work can first support lower local service costs. Exportable APIs can then pressure Claude pricing without customers buying that hardware. Only a reproducible hardware, software and support system across international deployments would establish a broader infrastructure substitute. These are conditional stages, with no assigned dates, market shares or probabilities. For Anthropic the first financial exposure is lower API revenue per accepted task, followed by utilization and the cash left after minimum compute payments.

4. Strategy: Capacity, Distribution and Constraints

Anthropic seeks to secure demand before compute becomes a commodity. The company develops proprietary applications while selling models through partners. Clouds provide customer access and capacity, but also receive part of Anthropic's spending and own competing offerings. The strategy works if growth in independent end demand pays for long-term compute contracts.

A cloud partner simultaneously supplies capacity, funds growth and influences customer access. AWS, Google Cloud and Microsoft Azure place Claude within the customer's existing procurement and control environment. For the partner, this creates compute demand; for Anthropic, it accelerates sales. But Google and Microsoft also promote their own AI offerings. Cooperation expands the market while preserving competition for the end interface. Anthropic's own Claude provides a direct customer relationship, while cloud marketplaces offer a route into organizations that a standalone application would reach more slowly. Cloud partnerships, Microsoft Copilot.

Different chip families provide compute choice but require engineering expertise in-house. Anthropic uses AWS Trainium, Google TPU and NVIDIA GPU; the prospectus also describes AMD GPU and a layer that allocates tasks across platforms. This allows the company to select available capacity for specific work and reduces dependence on one architecture. AWS nevertheless remains its primary cloud and training partner under the April announcement. The strategic benefit is supplier choice; it does not establish identical savings across all chips. AWS, Google/Broadcom, S-1/A, infrastructure.

Anthropic's cash capex is distinct from purchased cloud services and equipment or facility leases. In the first nine months of 2026, cash capex was $925mn, lease and equipment prepayments $2.0bn, and finance lease principal $88mn. Partners' infrastructure investment is not an additional Anthropic capex line. Direct procurement can change who owns or finances the equipment, but cannot erase already contracted cash minimums. Supplier and architecture flexibility must be demonstrated through porting cost, delivery, comparable workload economics and contractual rights. S-1/A.

Distribution must generate revenue before infrastructure becomes a fixed burden. Model tariffs and customer routing can change at the next release, while capacity is often procured 12–18 months before availability. Microsoft, Google and Amazon minimum-spend agreements extend into 2033 or 2036. Anthropic must keep refreshing its models and fill that capacity with paid work. The duration of an outcome premium is not disclosed, so the size of a purchased compute estate cannot establish the duration of shareholder returns.

Table 4 connects the five competitive forces with mechanisms that change profit.

Table 4: Competition and suppliers determine who captures compute savings

Force	Observable measure	Implication for profit
Rivalry	Three widely used models at $4 per basket	Advantage must be demonstrated through outcomes
Buyers	81% use multiple model families in a16z's sample	Tasks can be allocated among models
Suppliers	$69.8bn of commitments over the next 12 months	Utilization of paid capacity becomes critical
Substitutes	A compatible DeepSeek model for Claude Code	Application and model revenue can diverge

Force	Observable measure	Implication for profit
New entrants	Anthropic's $65.0bn funding round	Capital limits entry into frontier-model training

Sources: official Anthropic pricing, a16z, S-1/A, DeepSeek, Series H. The forces are not ranked without a measured customer response to price.

Compute and infrastructure obligations as of and after 30 September total $895.4bn, with $447.5bn due before the end of 2030 and most of the remainder between 2031 and 2045. The $213.8bn bucket with termination rights or pre-commencement conditions is not unconditional demand cancellation. Of $69.8bn due in the following twelve months, $4.3bn falls in that bucket. The remaining $65.5bn is outside it. These amounts describe future services and leases, not current cash debt, and they cannot be deducted again from a cash flow already reflecting the same payments. S-1/A, infrastructure.

The long-term table shows $320.0bn of determinable obligations, including non-cancelable purchases with more than a year remaining. It separately excludes $17.4bn of short-term arrangements that can be canceled within a year or expire within a year, without identifying how much belongs to each category. Uncommenced leases, later agreements and variable-capacity arrangements have different timing and conditions. Contract-by-contract cancellation costs, repricing and usable capacity are not fully disclosed. We do not call all $895.4bn fixed or all of the conditional bucket freely cancellable. S-1/A.

The disclosed minimum payments are contract specific. Microsoft requires at least $44.9bn from November 2026 to June 2033 regardless of usage, with relief for uncured breach or undelivered tranches. Google requires about $121.3bn from April 2026 to July 2033 and Amazon at least $110.0bn from May 2026 to April 2036. Their shortfalls can require true-ups, subject to contractual adjustments and exceptions. Lower demand or cheaper inference is not itself a general exit right. These arrangements are components of the obligation disclosure, not additions to its aggregate. S-1/A.

Political and legal conditions affect access to compute, data and customers. Their financial impact is visible, for example, in the Bartz settlement: a $1.5bn expense in 2025. The technological environment changes the price of the same task; the economic environment changes capital availability; and the social environment changes how much verification customers are willing to entrust to a machine. No separate premiums are assigned here for unmeasured risks. S-1/A, financial results and risks.

A common shareholder will receive an economic interest with limited control. Anthropic is a public benefit corporation, or PBC. The Long-Term Benefit Trust selects a majority of the board; the proposed Founder LLC share at the IPO will carry 50.1% of the votes on general matters without economic rights. The company has already contributed 846,840,960 Class N shares to the Public Benefit Fund—5% of fully diluted equity as of 30 June 2026, assuming full conversion of the Amazon notes. The board may issue up to a further 2% each time the share price subsequently rises 50% above the previous threshold, starting with the IPO price; this is discretionary authority, not a mandatory issuance. Investors must assess cash allocation in light of these powers, rather than assume they can compel distributions. Governance in S-1/A, Fund.

5. Financials and Owner Earnings

Anthropic has converted rapid growth into positive operating cash flow, but sustainable owner returns remain to be demonstrated. Recognized revenue is delivered service over a period, whereas an annualized pace is a point-in-time measure. May's announced pace exceeded $47.0bn; recognized revenue over the first nine months was $33.2bn. The S-1/A contains audited annual 2023–2025 financial statements and unaudited 2025 and 2026 nine-month figures. Series H, S-1/A.

Table 5 keeps annual and nine-month periods separate.

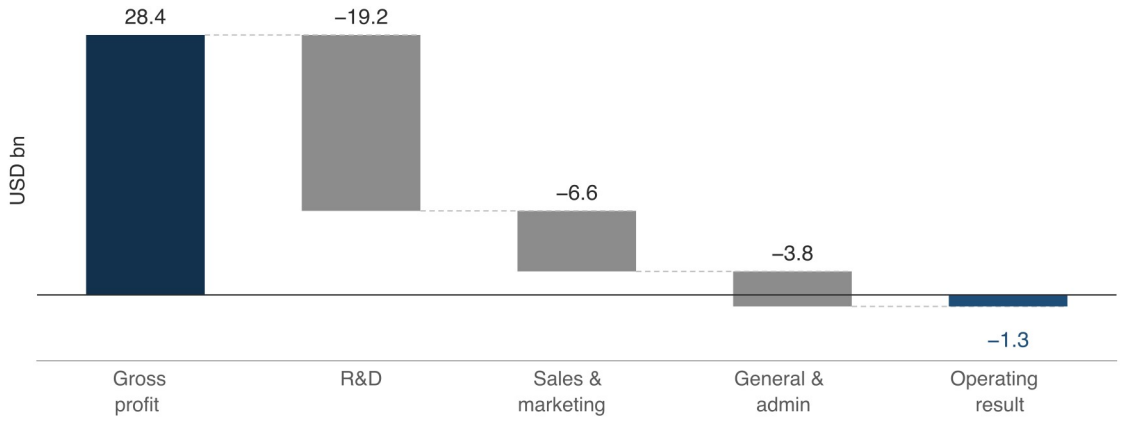
Table 5: Anthropic's financial history, $bn

Metric	2024	2025	9M 2025	9M 2026
Recognized revenue	0.4	4.6	2.5	33.2
Cost of revenue	0.8	2.8	1.8	4.8
Gross profit	−0.4	1.8	0.7	28.4
Gross margin	−94.8%	39.5%	29.4%	85.5%
Research and development	2.3	6.3	3.6	19.2
Sales and marketing	0.2	0.9	0.5	6.6
General and administrative	0.2	1.1	0.3	3.8
Bartz settlement	0.0	1.5	1.5	0.0
Operating income	−3.0	−8.1	−5.2	−1.3
Net income	−8.3	−42.0	−25.1	−104.0
Expense from remeasurement of financial instruments	5.4	34.3	20.1	94.3
Adjusted operating income[1]	−2.8	−4.9	−3.1	4.1
Operating cash flow	−5.7	−2.6	−2.0	6.6
Cash capital expenditure	<0.1	0.1	0.1	0.9
Prepayments for leases and equipment	0.0	0.1	0.1	2.0
Finance lease principal repayments	0.0	0.0	0.0	0.1
Cash remaining after the three preceding payments[2]	−5.7	−2.8	−2.2	3.6

Source: S-1/A, financial statements F-7/F-12, Note 8, reconciliation on pp. 158–159. Annual data are audited; interim data are unaudited. Gross profit, margin and the cash remainder are calculated from disclosures. [1] Company metric excluding, among other items, stock-based compensation and charitable commitments. [2] Operating cash flow less capital expenditure, prepayments and lease repayments; not owner earnings.

Our third conclusion is that profitability must survive the cost of repeated frontier refreshes. Note 8 allocates customer-serving workloads to cost of revenue and model development, research and internal usage mainly to R&D. The functional $19.2bn R&D expense overlaps with the $18.8bn compute-and-infrastructure expense classified by nature. Adding those two numbers would count costs twice. Functional expenses reconcile $28.4bn of gross profit to a $1.3bn operating loss. S-1/A, Note 8.

Chart 4: Gross profit of $28.4bn still leaves a $1.3bn operating loss

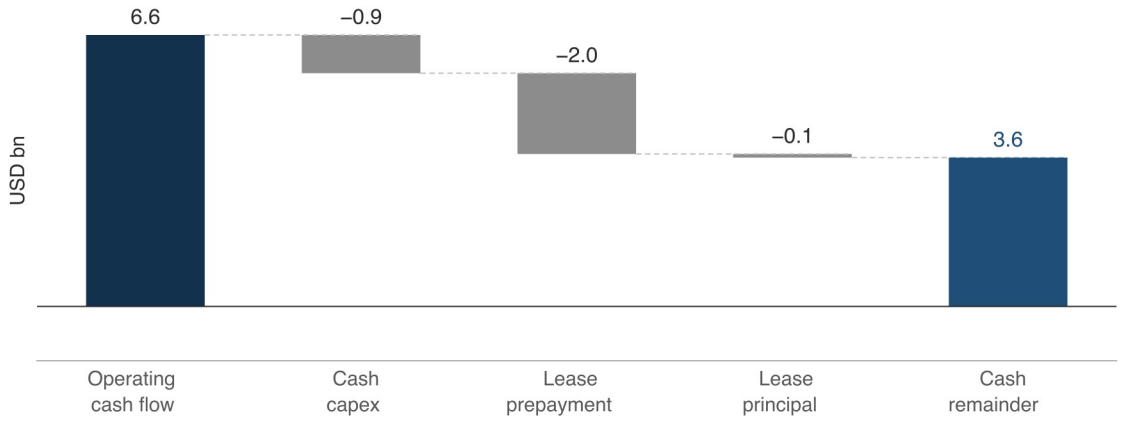


Source: S-1/A, first nine months of 2026. Research: $19.2bn; sales: $6.6bn; administrative expenses: $3.8bn.

The retained test is a sustainable shareholder cash margin of at least 25% after maintaining model quality and fully compensating employees. It is a reverse-valuation assumption, not an observed industry norm. Sustaining it for two years without losing comparable paid work would refute our concern. Holding other functional expense shares constant, R&D would have to decline from 57.9% to about 29.1% of revenue just to reach a 25% operating margin. That is a requirement, not a proposed research cut or a 25% cash-flow forecast.

The $104.0bn net loss includes $94.3bn of fair-value remeasurement of financial instruments. That charge is not an equal cash burn and the related conversion rights cannot be charged a second time as both full debt and full dilution. Operating cash flow already includes cash research costs. Assessing how much future R&D maintains competitiveness informs the forward margin, but does not authorize subtracting the same research again from historical operating cash flow. Stock compensation remains a cost to owners through cash compensation or dilution. S-1/A.

Chart 5: $3.6bn remains after capital and lease payments



Source: S-1/A, cash flows for the first nine months of 2026. Deductions are $0.9bn, $2.0bn and $0.1bn, respectively. The remainder is not normalized owner earnings.

The $3.6bn cash remainder is operating cash flow less cash capex, lease prepayments and finance lease principal. It is not normalized distributable cash. Receivables, payables, advance payments and the timing of capacity delivery can move it materially. More efficient software improves this remainder only when it avoids a payment,

replaces expensive short-term compute, or earns additional net revenue on already-paid capacity. A price cut can absorb the saving; unused committed capacity can leave payments unchanged. S-1/A.

Business quality across a full cycle has not yet been established: there is no sufficient observation window for returns against the cost of capital; the maintenance requirement has not been separated; the advantage in the worst year of that window has not been measured; and receipt by minority shareholders of at least one-third of earnings has not been confirmed. This is missing evidence of quality, not a conclusion that the business is poor.

6. Funding Rounds and Reverse Valuation

The latest announced round demonstrates historical access to capital, not a 6 November share close. Announced valuation rose from $61.5bn at Series E to $965.0bn at Series H. Preferred and common rights differ. Anthropic remains private for this review, so there is no regular-session 6 November public close, current public market capitalization, public return-to-target calculation or Nasdaq-derived upside. The round reference is neither a current common transaction nor a new public target.

Table 6 shows the announced terms of each round separately.

Table 6: Latest announced funding rounds, $bn

Announcement	Round	Announced amount	Post-money valuation
3 March 2025	E	3.5	61.5
2 September 2025	F	13.0	183.0
12 February 2026	G	30.0	380.0
28 May 2026	H	65.0	965.0

Sources: Series E, F, G, H. Announced amounts do not add up to cash received.

A primary-source inconsistency remains: the H announcement includes $15.0bn of earlier partner commitments; Note 9 labels Google's and Amazon's May investments H, while Note 10 and Recent Sales separately identify $15.0bn of G-2. We retain that inconsistency. Financial statements show $110.4bn of net preferred proceeds for the first nine months. Shareholder resales are not company financing; the October 2025 Series B resale generated zero net proceeds for Anthropic. S-1/A.

Cash and short-term investments were $131.8bn at 30 September. The $14.8bn available credit and up to $77.0bn of undrawn partner note capacity are conditional funding sources, not cash. Amazon notes have $6.2bn face principal and a $130.5bn fair value. Expected share settlement makes conversion relevant to ownership; maturity cash or set-off risk remains under the terms. Neither the fair-value charge nor note fair value is a second cash outflow. Final post-IPO fully diluted shares remain undetermined. S-1/A, balance sheet, Notes 9–10.

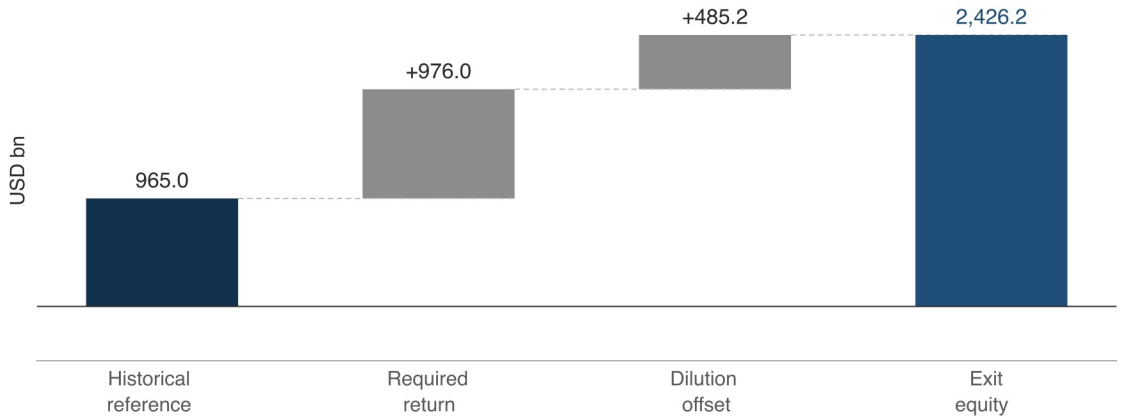
The $131.8bn cash-and-investment balance exceeds the disclosed next-twelve-month infrastructure obligations by $62.0bn before all other receipts, outflows and restrictions. This static comparison is not a runway forecast. It excludes the timing of collections, payroll, taxes, research payments outside that obligation perimeter and future additions to capacity. If lower task prices coincide with underutilization and mandatory payments, incremental funding needs can rise. New equity or conversion can then reduce each existing holder's share of the eventual cash flow. S-1/A.

A reverse illustration separates entry requirements from a forecast. The historical Series H reference is $965.0bn; the hypothetical entry date is 5 November 2026 and exit 5 November 2031. A 15% annual return, no interim distributions and 20% loss of the original ownership require $2,426.2bn of exit equity. At 25 times owner cash

flow and a 25% cash margin, annual revenue must reach $388.2bn. None of these assumptions is a current transaction or measured cost of equity.

Third-quarter revenue of $17.0bn gives $67.8bn when multiplied by four. This is a simple annualization, not contracted recurring revenue. The $388.2bn reverse requirement is 5.72 times the annualized quarter using unrounded values, equivalent to about 41.8% annual growth over five years. The same result also requires the retained 25% owner cash margin after frontier upkeep and reinvestment. Revenue growth cannot substitute for that margin. S-1/A, quarterly results.

Chart 6: A 15% annual return requires $2,426.2bn of equity value in 5 years



Source for the starting reference: Series H, 28 May 2026. The rest is the author's conditional calculation for 5 November 2026–5 November 2031; not a valuation forecast.

Table 7 translates the owner's required outcome into company cash flow and revenue.

Table 7: Requirements of an illustrative equity owner on 5 November 2031

Measure	Value
Required value before dilution: 965×1.15^5	$1,941.0bn
Total equity value after 20% dilution	$2,426.2bn
Annual shareholder cash flow at a price of 25 times that flow	$97.0bn
Revenue at a 25% cash margin	$388.2bn
Revenue at a 20% cash margin	$485.2bn
Revenue if entry reference rises 10% to $1,061.5bn	$427.0bn

The requirement compounds the historical entry reference at the chosen return, divides by the retained ownership share, and then by the equity-to-cash-flow multiple and cash margin. The complete four-step calculation is repeated for every boundary in Charts 7 and 8. This is total equity arithmetic, with no reliable post-IPO per-share denominator. Future cash and debt are not separately added to an equity/cash-flow multiple. The common rate inputs are 6 November Rf 4.86% and the retained ERP convention of 4.5%, with 5.0% a separate sensitivity. Without a supported issuer beta or private premium they do not replace the diagnostic 15% hurdle. US Treasury, 6 November 2026.

Chart 8's Entry reference +10% raises entry from $965.0bn to $1,061.5bn and required revenue from $388.2bn to $427.0bn (up $38.8bn), with the return, horizon, ownership loss, multiple and margin unchanged.

A 20% cash margin instead of 25% raises required annual revenue from $388.2bn to $485.2bn. Raising ownership loss from 20% to 30% alone requires $443.6bn. Combining both changes requires $554.6bn, an increase of $166.4bn from the exact base. These effects are not added mechanically. They show how weaker retained economics and financing dilution compound. They imply neither an IPO return nor fair value per common share.

The valuation question is whether sufficiently many paid outcomes can earn that residual margin for long enough. Faster inference can help utilization and avoid expensive marginal purchases, while repeated model refreshes consume resources and lower API prices return benefits to buyers. A shorter premium would weaken the exit cash multiple: the 25x-to-20x boundary also requires $485.2bn of revenue at the base cash margin. This does not measure a particular number of years of advantage. The six-year exit boundary requires $446.4bn because the owner waits longer for the same annual return; it is not an estimate of a longer moat.

Chart 7: A 20% cash margin and 30% ownership loss require $554.6bn of annual revenue

Required annual revenue at exit | 5 Nov 2031 exit horizon

Ownership loss	Cash margin after upkeep: 25% cash margin	20% cash margin
20% ownership lost	388.2	485.2
30% ownership lost	443.6	554.6

Joint boundary = 166.4 additional annual revenue, USD bn

Source: conditional reverse-valuation assumptions for an illustrative 5 November 2026 entry. USD bn at 5 November 2031. Entry $965.0bn, return 15%, exit multiple 25x, no interim distributions. Required revenue, not V0 or TP12m.

Chart 8: Individual assumption changes require $427.0bn to $485.2bn of annual revenue

Only the named input changes | 6-year exit: 5 Nov 2032

Changed input	Required annual revenue at exit: Retained base	Saved boundary
Return 15% to 18%	388.2 Δ +0.0	441.5 Δ +53.3
Ownership lost 20% to 30%	388.2 Δ +0.0	443.6 Δ +55.5
Cash margin 25% to 20%	388.2 Δ +0.0	485.2 Δ +97.0
Exit multiple 25x to 20x	388.2 Δ +0.0	485.2 Δ +97.0
Exit after 5 to 6 years	388.2 Δ +0.0	446.4 Δ +58.2
Entry reference +10%	388.2 Δ +0.0	427.0 Δ +38.8

5 pp lower cash margin = 97.0 more annual revenue, USD bn

Source: conditional reverse-valuation assumptions. Each row changes only its named input. Entry: 5 November 2026. Exit: 5 November 2031, or 5 November 2032 for the six-year row. Amounts and deltas are USD bn. Utilization and frontier-refresh cash needs remain unmeasured.

7. Scenarios, Observations and Risks

The decisive scenario is retaining customers as models become cheaper and paying to maintain quality. There is no central forecast, scenario probabilities or investment rating here. The business is a potential compounder with a long reinvestment runway; compounding of value per share has not yet been established. The observations below could change our position; they are not promised financial results.

Table 8 sets out future observations and conditions for changing our position.

Table 8: What must confirm or refute the investment story

Decision point	Next observation and timing	Threshold for changing our position
Quality premium	Identical work tasks by 31 March 2027	A cheaper model delivers the same accepted outcome with less than 10 seconds of additional verification in the specified example
Pricing power	Two quarters after an actual price increase	A 10% price rise without a loss of the share of comparable paid tasks or offsetting discounts
Owner earnings	2027–2028 reports	25% cash margin for two years after frontier refresh, all mandatory payments and the full compensation cost, without deducting historical R&D twice
Capacity	The next 12 months of disclosures	Paid use covers minimum capacity payments; distinguish timing relief, supplier breach and pre-commencement conditions from demand cancellation
Ability to execute an investment	Final IPO prospectus	Price, share count, conversion and new issuance allow common share returns to be calculated

The first three thresholds are specified in advance for subsequent testing; results have not been obtained. Market consensus on these measures has not been established, so the research does not claim a demonstrated divergence from the market.

The economic path runs from accepted paid tasks and their realized price to revenue, then through utilization, serving costs, research refresh, other operating costs, cash payments and financing to cash per owner. For a fixed minimum payment M and paid output Q, the payment per output is M/Q: filling idle paid capacity reduces this unit amount but does not itself reduce M. Retained efficiency, price response, paid-task utilization and payment schedules are not disclosed well enough to assign a central cash-flow path. They remain unknowns, rather than invented savings percentages or financing probabilities.

Permanent loss arises if the economic life of the premium is shorter than the payments it must finance. A sufficiently good cheaper model takes tasks or forces lower prices, while compute minimums continue. Research cannot be cut repeatedly without testing the quality premium. Underutilization and weaker cash generation can then require financing on worse terms and reduce existing ownership. The joint retained boundary in Section 6 raises required revenue to $554.6bn. It measures conditional risk size; no causal link between a technical publication and a market fall has been established. S-1/A.

Our position is to understand and monitor the business, with no investment rating or price target before executable terms. The next tests are accepted-outcome cost, the response of comparable paid work to actual price changes, cash after all recurring research and mandatory capacity payments, and the ownership left after financing. Final IPO pricing and fully diluted shares would make a public shareholder return calculable. Strong growth and access to capital support the business, but do not by themselves establish a reason to buy at the historical round mark.

8. Sources

1. **Financial statements and prospectus:** Anthropic S-1/A, 2 November 2026. Principal source for financial history, cost of revenue, commitments, capital and governance; initial S-1, 26 October 2026. The IPO had not completed as of the cut-off date.

2. **Funding rounds:** official announcements for E, F, G, H. Announced terms are kept separate from receipts during the reporting period.

3. **Products and pricing:** Opus, Sonnet, Claude, OpenAI Sol, Google, DeepSeek. Pricing snapshot as of 5 November 2026; suppliers' actual costs are not inferred from price lists.

4. **Market measurements:** Menlo, 3 November 2026, a16z, 30 January 2026, Ramp AI Index. Populations and units of measurement differ; Menlo is an investor in Anthropic, and a16z discloses an investment in OpenAI.

5. **Outcome verification:** METR, 2025 experiment, Claude Code research. Older experiments do not replace measurement of current products.

6. **Shareholder rights:** Long-Term Benefit Trust, Public Benefit Fund and the relevant sections of S-1/A.

Report dated 7 November 2026. Financial and market information has a 5 November cut-off unless specifically stated. S-1/A of 2 November is checked for the audit scope, functional expenses and contractual terms. No completed IPO or public share close is inferred. Human acceptance and release sign-off have not been recorded.

Appendix Entry requirements and evidence

Paired analytical workbook: Anthropic — модель.xlsx; reverse-entry assumptions dated 5 November 2026, reviewed 7 November. The $965.0bn entry mark comes from the historical 28 May Series H announcement; no current common quote, public rating, V_0 per share or TP12m is issued.

The full reverse chain starts with an investor's required proceeds. $965.0bn × 1.15^5 = $1,941.0bn; retaining 80% of the initial ownership requires total exit equity of $2,426.2bn. Dividing by the 25x equity/cash-flow multiple requires $97.0bn of annual distributable cash. Dividing by the 25% margin requires $388.2bn of annual revenue. The corresponding entry check runs the chain backwards: entry equity = revenue × cash margin × cash-flow multiple × retained ownership ÷ $(1 + \text{return})^5$. Substituting the exact base reproduces $965.0bn. This is equity-to-equity arithmetic: adding cash or subtracting debt again would change the stated method.

The economic scale check uses the same revenue definition. Required $388.2bn is 5.72 times the annualised Q3 recognised revenue and requires about 41.8% annual growth over five years, while preserving the owner cash margin. At a 20% margin the requirement becomes $485.2bn. These comparisons test what the historical entry demands; revenue run rates, API spending shares and private preferred rounds do not supply comparable public trading multiples. No weighting of DCF, rounds or a peer multiple is introduced. Charts 7 and 8 retain the joint 2×2 and individual full-recalculation boundaries.

Reported values come from the 2 November S-1/A and dated funding announcements; financial history remains at the 5 November cutoff. Calculated values include the reverse chain and the $427.0bn entry-plus-10% result. Estimated inputs are the 15% hurdle, 20% ownership loss, 25x multiple and 25% normalised cash margin. Inferred claims concern whether task quality can defend paid demand and retained cash; they are not observed margins or consensus. Final IPO rights and fully diluted shares remain undetermined.

Table 9: Cash margin and financing can raise the required scale materially

Material limitation	Conditional impact or exposure	Next observation and responsible work
Normalised owner cash after frontier upkeep	25% to 20% margin raises required revenue from $388.2bn to $485.2bn; upkeep/unit cost unmeasured	Issuer analyst: 2027–2028 cash reports, recurring research, working capital and all capacity payments
Financing and retained ownership	20% to 30% ownership loss alone requires $443.6bn; jointly with 20% margin requires $554.6bn	Capital analyst: final IPO prospectus, note conversion and subsequent issuance
Capacity utilisation and cancellation rights	$69.8bn obligations over the next 12 months; no measured standalone valuation charge	Issuer analyst: paid use, commencement, minimum-payment and relief terms at each disclosure
Executable entry and common-share rights	No reliable post-IPO per-share denominator; investment return at actual terms not calculable	Capital analyst: final pricing, fully diluted capital table and share-class terms before any rating

Table 9 records exposures and requirements, not probability-weighted loss estimates. Unmeasured effects are not assigned zero. No public TSR signal or active recommendation follows from the historical private mark.

Rounding: calculations retain full precision. Percentages, ratios and totals are calculated before display rounding; shown rows may not sum exactly. USD bn is displayed to one decimal; small material amounts are identified in USD m.