The Anthropic Three-Statement Model

What must you believe for $2 trillion (v2.2, audited)

Valuation date 31 Dec 2026 (pre-IPO) | Basis: fully diluted equity value post-raise | Explicit forecast window 2026-2035 | All figures $bn unless noted

Model workbook Anthropic_3S_Valuation_Model.xlsx (11 sheets, including the cash-flow / balance-sheet articulation); engine engine.py, every assumption explicit and editable

Data as of 2026-11-03; roadshow figures are unaudited on a company-selected basis - every input is treated as [EST]. Author: Stanley Choi

0 The one-line conclusion

$2 trillion is almost exactly a demand that the entire roadshow curve come true: it corresponds to 2030 revenue of $935bn - 84% of the roadshow's own $1,120bn projection. The evidence alone (the measured deceleration curve plus the commitment book) supports about **$287bn**; probability-weighted, about **$893bn**. The $1.9tn ask implies assigning roughly a **34%** probability to the roadshow path - 6.6x the value of the evidence.

1 First, reconcile the disclosures: three “contradictions” are one economics

The roadshow claims at once: an 85% adjusted gross margin, a compute margin above 40% ($1.70 of revenue per $1.00 of compute), and 70-75% of compute going to training. Rebuilt from actual Q3 revenue of $16.956bn:

Item	Amount ($bn)	Share of revenue
Revenue	16.956	100%
Total compute cost (= 16.956 / 1.72)	9.863	58%
- training compute (75%, in R&D;)	7.431	44%
- inference compute (in COGS)	2.432	14%
Adjusted COGS (implied by the 85% gross margin)	2.432	14%
Adjusted operating income (11%)	1.87	11%

Inference compute (14%) roughly equals adjusted COGS (14%) - the three numbers are fully self-consistent. The price is that the “41% compute margin” is really an all-in gross margin, and the “85% gross margin” holds only because seven-tenths of the compute is booked as R&D.; If the model race keeps training spend above 40% of revenue, this margin structure never converts into operating leverage.

2 The touchstone, quantified: what commitment rigidity does to the margin

The $895.4bn of commitments ($447.5bn before 2030), about 87% non-cancellable, is laid out as an annual payment curve on a capacity ramp; any excess of the scheduled payment over the scenario's compute usage (forced over-consumption) enters operating income directly. Under stress (S3):

$bn	2026	2027	2028	2029	2030	2031	2032
Forced overpay (locked payments above usage)	0	40.0	60.3	76.5	88.0	56.3	51.6
Adjusted operating income (incl. forced overpay)	5.9	-35.4	-57.1	-75.1	-88.0	-57.6	-52.2
True FCFF (adding SBC and financing)	-1.7	-44.1	-66.6	-87.6	-100.8	-70.2	-64.5

In 2030, a company with an 11% book operating margin becomes a cash furnace running at -$100.8bn of FCFF, more than the year's revenue ($67bn), under commitment rigidity; book equity turns negative in 2028 (technical insolvency) and operations continue thereafter on injections and 8% debt. “Positive FCF” and “the $518bn book” cannot be read separately - this is that statement in arithmetic. On the consensus path 2026 FCFF is -$0.1bn - break-even at best, with no buffer; the comfort in “adjusted profitability” comes entirely from leaving SBC and the

commitment payments outside the basis.

3 The three paths in full

	S1 roadshow as pitched	S2 consensus (measured)	S3 circular stress
2030 revenue	$1,120bn (company projection)	$175bn (S-curve)	$67bn (open-source deflation)
2035 revenue	$2,486bn	$314bn	$71bn
2035 FCFF	~$620bn	$41bn	-$26bn (book insolvent)
Equity value (WACC 10.5%)	$5.1tn	$287bn	$0 (renegotiation floor)
Implied 2026 ARR multiple	74x	3x	-

The S2 curve is anchored on measurement: the ARR monthly compounding rate of 32% from February to July 2026 falls to 2% from August (distillation crackdowns, the OpenAI fightback, Meta and Microsoft cutting internal Claude spend, price war), with a mild 2027 recovery and deceleration thereafter (+49% to +40% to +32% to +25% ... to +7%), reaching 16% of the roadshow number by 2030. S1's 2031-35 is a decelerating extension.

4 Reverse DCF: what $2tn / $1.5tn / $1tn each require

Back-solved from the S2 consensus path (WACC 10.5%, terminal growth 3%):

To reach the target, need...	$2.0T	$1.5T	$1.0T
2030 revenue (consensus $175bn)	5.34x, or $935bn	3.99x, or $699bn	2.64x, or $462bn
or a WACC of	4.6%	5.1%	6.2%
or operating margin up by	+40pp a year (hits the 60% cap)	+40pp (cap again)	+27pp (49% by 2030)
or flat 2027-35 FCFF of	$174bn a year	$129bn a year	$85bn a year
or the roadshow path realised at	39% (2030 = $438bn)	29% ($329bn)	20% ($224bn)

Reading: **$1.0tn needs 2030 revenue of $462bn (2.6x consensus), a 6.2% WACC, or a 49% operating margin** - none absurd, but each more optimistic than the measured extrapolation: it is the defensible ceiling. $1.5tn begins to require miracles of margin or discounting (the margin solver saturates at its 60% cap). **$2.0tn corresponds to $935bn of 2030 revenue - essentially the roadshow itself** (84% of it) - a price that assumes the roadshow curve is eighty-percent true, which the measured deceleration says it is not.

5 Best estimate

Scenario	Equity value	Subjective probability	Basis
S1 roadshow roughly holds	$5,093bn	15%	needs 2030 revenue of $1.1tn = 16x 2026, with the Mag7 capex pool and external paying capacity both unlimited
S2 consensus deceleration	$287bn	45%	the measured deceleration curve plus still-expanding enterprise AI spend plus locked-capacity advantage
S3 circular stress	$0	40%	open-source substitution plus price deflation plus customer concentration plus commitment rigidity; the credit data (Oracle CDS 261bp, private-credit gating) supports this weight
Probability-weighted	~ $893bn		the $1.8-2tn ask implies ~34% probability on the roadshow path

Best estimate: **$290-450bn** (anchored on S2, widened both ways), with the $893bn probability-weighted value as the equilibrium price when sentiment allows the roadshow a one-in-three weight. The evidence is worth $287bn; the $1.9tn ask is 6.6x that - not fundamental pricing, but the probability-selling of a growth curve.

6 Evidence grading of the assumptions

Grade	Items

A - Measured	Q1-Q3 2026 revenue (4.731 / 11.555 / 16.956); FY25 revenue 4.585 / operating loss 8.060 / net loss 41.965 / compute 7.326; the ARR curve (monthly growth 32% falling to 2%); 895.4 of total commitments (447.5 before 2030)
B - Roadshow, unaudited	85% gross margin; 11% operating margin; $1.70/$1.00; NRR 900%; 80% non-cancellable; the 2030 projection of $1,120bn
C - Model inference	training share 72.5%; SBC at 9% of revenue; $30bn IPO raise; the 2031-35 extension; WACC 10.5%; terminal g 3%; compute intensity 50-62%

7 Funding and plug policy (explicit)

The three statements use **equity injection as the plug account**: each year the cash-flow statement first sizes the gap between cash (with a minimum-cash floor of 15% of current-year revenue) and this year's FCFF; 50% of the gap is funded with debt up to the scenario cap (S1 $300bn / S2 $60bn / S3 $35bn - the S3 cap reflects private-credit gating), and the remainder is booked as an equity injection. Debt costs 8% (2026 AI-infrastructure private-credit pricing: CoreWeave A3 at SOFR+225bp, junk at 9-15%); the cash tax rate is 21% (federal; state folded into NOL conservatism); injections are booked at book/par with no down-round discount modelled - the cost is that under stress book equity turns negative in 2028 (technical insolvency shown honestly), and the valuation separately floors S3 equity at zero. Plug usage: S1 never triggered; S2 $24bn of debt and $24bn of injections over the decade; S3 $419bn of injections plus $35bn of debt (at the cap).

8 Quality gates (v2.2 audit results)

Gate	Result
Balance sheet articulation (A = L + E, every scenario and year)	PASS (max deviation 4.6e-13)
Funding plug policy made explicit	PASS (minimum cash 15% of revenue; gap funded debt-first; plug = equity injection, visible line by line in the CF_BS sheet)
Cash-flow to income articulation (CFO+CFI+CFD = delta cash)	PASS (guaranteed by identity)
Workbook formulas vs engine, cell by cell (3 scenarios x 10 years)	PASS (formula-level audit 2026-10-17: the forced-overpay omission found and fixed; zero mismatches)
Assumptions separated from calculations	PASS (standalone Assumptions sheet; the calculation layer references it by formula)
Basis consistency	PASS (adjusted basis throughout = ex-SBC; SBC charged to FCFF at cash value and credited to APIC)
Terminal-value share disclosed	68% in both S2 and S1 - the result leans on the terminal assumption; a WACC 8-13% sensitivity is provided
v1 to v2.1 audit corrections	removed the unsupported D&A; add-back (which inflated FCFF by $5-12bn a year); unified the working-capital basis; booked injections into equity; calibrated debt cost and tax rate to 2026 market reality

9 What the model does not handle (honest disclosure)

The dilution path: SBC is charged to FCFF at cash-equivalent value and credited to APIC, but post-IPO share issuance (warrants, acquisitions, later rounds) is not modelled line by line. SPV/lease bankruptcy remoteness: vendor financing is disclosed only; no second-stage contagion from funder repurchase or recourse is modelled. Cross-audit: the Anthropic and OpenAI commitments are each other's demand assumptions (the consolidation logic); this model is single-entity. All roadshow inputs are unaudited, and the adjustments in “adjusted” are company-selected (not enumerated officially).

10 Disclaimer

This model and this report are for financial analysis and teaching purposes only. They are not investment advice, a price target or a rating. All forward-looking figures are estimates based on explicit assumptions [EST]; actual results may differ materially as company disclosures, market conditions and the assumptions themselves evolve.

Model: engine.py (pure stdlib, deterministic, with a year-by-year balance assertion) plus build_xlsx.py; the workbook formula chains reconcile with the engine cell by cell. English edition, 2026-10-17. Stanley Choi - Queensland University of Technology