The OpenAI Three-Statement Model

What must you believe for $1 trillion

Valuation date 31 Dec 2026 (IPO expected 2027; the $60bn raise is booked in 2027) | Basis: fully diluted equity value post-raise | Explicit forecast window 2026-2035 | All figures $bn unless noted

Model workbook OpenAI_3S_Valuation_Model.xlsx (11 sheets, including the cash-flow / balance-sheet articulation); engine engine.py, every assumption explicit and editable

Data as of 2026-10-17 (site transaction ledger as of 10-11); the internal plan and roadshow figures are unaudited - every input is treated as [EST]. Author: Stanley Choi

---

0 The one-line conclusion

Pointing the same machinery as Appendix A at OpenAI produces a more extreme result: on the consensus path the equity is consumed before the terminal year by the commitment book (~$1,360bn, roughly 34x its ARR) - the value is floored at 0. The only path worth anything is the company's own 2026-30 internal plan, and it is worth $849bn - almost exactly the $852bn post-money of the March-2026 private round. The private market has already priced the plan as fully true. Probability-weighted, the equity is about **$170bn**; **a $1.0tn IPO is not reachable even at 100% probability on the roadshow path** (the roadshow itself is worth $849bn).

1 Reconciling the bases: apparent loss, cash burn and ARR

Media basis: a 2025 apparent loss of $38.5bn, but an adjusted cash burn of about $8bn - the gap is SBC, remeasurements and non-cash items. For 2026: full-year revenue of about $36bn (Q1-Q3: 7.6 / 8.9 / 10.5), an announced ARR of $40bn in August, against 2026Q2 annualised booked revenue of only $26.8bn - an $13bn gap between announcement and booking. The model runs on the adjusted basis and charges SBC separately, at 12-20% of revenue (higher than Anthropic's 6-12%).

2 The touchstone, quantified: the $1,360bn commitment book

The commitment composition: Oracle $300bn over five years (Stargate), Microsoft Azure $250bn, Nvidia $100bn (a 10GW LOI with $30bn deployed plus a $105bn lease guarantee), AWS $138bn, AMD 6GW at roughly $350bn [EST] (with a warrant for about 10%% of the company), Broadcom 10GW at roughly $200bn [EST], CoreWeave $22.4bn - about $1,360bn in total, roughly 34x the announced $40bn ARR (Anthropic's book is 5.2x its ARR). Laying the schedule out on a capacity-ramp curve, any year in which the scheduled payment exceeds the scenario's usage enters the P&L as forced over-consumption:

$bn	2026	2027	2028	2029	2030	2031	2032
Commitment schedule	30	95	150	185	200	185	175
S2 forced overpay	3	54	90	102	91	50	18
S2 true FCFF	-20	-76	-117	-133	-127	-85	-46
S3 forced overpay	0	57	108	140	153	140	131

On the consensus path, a company with $145bn of 2030 revenue is forced to overpay by $91bn - a book operating margin that has just turned (-2%) is swallowed whole; FCFF is negative throughout and book equity turns negative from 2028. Under stress the 2030 forced overpay is $153bn, about three times that year's revenue. That is the arithmetic of a book at roughly 34x ARR: revenue falling short of plan does not reduce the payments, it converts them into losses.

3 The three paths in full

	S1 internal plan as written	S2 consensus (measured)	S3 circular stress
2030 revenue	$350bn (FT-cited internal plan)	$145bn (S-curve)	$52bn (open-source deflation)

2035 revenue	$830bn	$272bn	$58bn
2026-30 cumulative FCFF	-$273bn (vs the internal -$278bn)	-$473bn	-$590bn
2035 FCFF	$138bn	-$9bn	-$58bn
Equity value (WACC 10.5%)	$849bn	$0 (renegotiation floor)	$0 (renegotiation floor)
Implied Aug-26 ARR multiple	19x	-	-

The S1 calibration anchor: the engine's 2026-30 cumulative FCFF of -$273bn sits within 2% of the company's own internal forecast of -$278bn (FT-cited) - the model reproduces the burn the plan itself admits to. The S2 revenue curve is anchored on the gap between the announced ARR and annualised booked revenue, reaching 41% of the internal plan by 2030. S1's EV is almost entirely terminal value: a decade of burn consumes the present value of the interim cash flows.

4 Reverse DCF: what each target requires - and which levers are broken

Back-solved from the S2 consensus path (WACC 10.5%, terminal growth 3%). Unlike Anthropic, three of the four levers are broken here:

To reach the target, need...	$1.5T	$1.0T	$852bn (private round)
2030 revenue	not reachable - on a negative-margin path, scaling revenue only scales the loss (the solver hits its 30x cap and still returns $0)	not reachable (same)	not reachable (same)
or a WACC of	below the 4% floor and still short	below the floor	below the floor
or adjusted op. margin up by	+40pp a year (the cap) still not enough - in 2030 the $200bn schedule exceeds usage even at a 38% margin; forced overpay absorbs the uplift	same	same
or flat 2027-35 FCFF of	$127bn a year	$82bn a year	$69bn a year
or the internal plan realised at	1.67x (2030 = $586bn)	1.21x (2030 = $422bn)	1.01x (2030 = $352bn)

Reading: **no single lever on the consensus path reaches even $1.0tn** - the growth multiplier and the margin uplift are both swallowed by commitment rigidity, and discounting cannot rescue negative flows. The only arithmetically viable route is flat 2027-35 FCFF of $82bn a year (the consensus 2035 FCFF is -$9bn), or the internal plan over-delivering. This is not Anthropic's structure, where $1.0tn was a defensible ceiling; OpenAI's consensus path has no ceiling in the model, only a floor (zero).

5 Best estimate

Scenario	Equity value	Subjective probability	Basis
S1 internal plan as written	$849bn	20%	needs 2030 revenue of $350bn = 10x 2026; the engine's cumulative burn is calibrated to the internal forecast; the private round priced exactly this
S2 consensus deceleration	$0	45%	the book is 34x ARR; any revenue miss becomes forced overpay; book equity negative from 2028
S3 circular stress	$0	35%	open-source substitution plus price deflation plus customer concentration; Oracle CDS and private-credit gating (ch.08) support the weight
Probability-weighted	~ $170bn		reference: March-2026 private round $852bn; media-speculated IPO $1tn+

Best estimate: probability-weighted equity of $170bn against the $852bn last private round - **the private market bought the roadshow path itself**. An IPO priced above $1tn asks the public market to pay more than the private round, and more than any probability mix the model can form.

6 Evidence grading of the assumptions

Grade	Items
A - Measured	FY25 revenue 13.0 / apparent loss 38.5 / adjusted cash burn 8.0; 2026 revenue ~36 (Q1-Q3: 7.6/8.9/10.5); ARR 40 vs 26.8 annualised booked; the March round of $122bn at $852bn post-money; SoftBank's $40bn bridge ($25.9bn repaid); Nvidia $30bn deployed plus a $105bn guarantee
B - Internal / leaked filings (unaudited)	the FT-cited internal forecast (2026 $36bn to 2030 $350bn, cumulative FCF -$278bn); the commitment composition of ~$1,360bn
C - Model inference	training share 60-70%; SBC at 12-20% of revenue; $60bn IPO raise (2027); end-2026 cash $60bn; the 2031-35 extension; WACC 10.5% (deliberately the same as Appendix A); terminal g 3%; compute intensity 60-90%; 60% non-cancellable

7 Funding and plug policy (explicit)

The same six-line policy as Appendix A: **equity injection is the plug account** - each year the cash-flow statement sizes the gap between cash (with a minimum 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 $500bn / S2 $150bn / S3 $60bn) and the remainder is booked as an equity injection at book value (no down-round discount modelled). Debt costs 8%; the cash tax rate is 21% (accumulated NOLs of about $60bn shelter early profits); the $60bn IPO raise enters 2027 cash and equity. Plug usage: S1 debt $77bn (the 2027-29 ramp); S2 debt $150bn (at the cap) plus $269bn of injections; S3 debt $60bn (at the cap) plus $932bn of injections. S2 book equity is negative from 2028 - insolvency is shown honestly, and the valuation separately floors S2/S3 at zero.

8 Quality gates

Gate	Result
Balance sheet articulation (A = L + E, every scenario and year)	PASS (max deviation 2e-13)
Cash-flow to income articulation (CFO+CFI+CFD = delta cash)	PASS (guaranteed by identity)
S1 calibration: engine 2026-30 cumulative FCFF vs the internal forecast	PASS (-273 vs -278, a 2% gap)
Workbook formulas vs engine, cell by cell (formula-level audit)	PASS (9 formula lines x 3 scenarios x 10 years reconcile; the forced-overpay omission found and fixed 2026-10-17)
Assumptions separated from calculations	PASS (standalone Assumptions sheet, evidence-graded; 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	S1's EV is almost 100% terminal - more extreme than Anthropic (68%); a WACC 8-13% sensitivity is provided
Comparability with Appendix A	same engine, same WACC/g, same plug policy - every difference is a company-level assumption

9 What the model does not handle (honest disclosure)

The post-recap structure (Microsoft roughly 27%, SoftBank roughly 13%) is approximated by cumulative investment; the dilution from the AMD warrant (~10%) and the Nvidia/CoreWeave stakes enters only through the fully-diluted basis, not milestone by milestone. SPV/lease bankruptcy remoteness: the $105bn lease guarantee and the Broadcom custom-chip financing are disclosed only; no second-stage recourse contagion is modelled. Cross-audit: the Anthropic and OpenAI commitments are each other's demand assumptions (the consolidation logic of Chapter 01); this model is single-entity. Advertising and consumer revenue are not modelled separately.

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