

Anthropic Before the Prospectus

Working Paper 1: An Analogical Risk Inventory

Robert G. Eccles and Shivaram Rajgopal

Information cutoff: September 26, 2026. Working paper accompanying our Heavy Lifting essay on Anthropic's forthcoming prospectus. Nothing in this analysis was non-public as of the cutoff; every factual claim is traceable to the source register in Appendix A, retrieved on September 26, 2026. On that date Anthropic's public Form S-1 had not been filed. This version corrects the treatment of US bank model-risk guidance, which was revised in April 2026.

Provenance and contamination notice. This draft was produced by Claude, a model made by Anthropic, working under the authors' direction. Claude holds no non-public information about Anthropic. It is not a neutral analyst of its own developer, and it was trained on a large volume of Anthropic's published self-description, so passages that read as analysis may be recall. We use four conventions in the text. [S##] cites the source register. **[unsourced recall]** marks any claim we could not trace to a retrieved source. **[our estimate]** marks a figure we derived, with the derivation shown. **[company statement only]** marks a claim that rests solely on Anthropic's self-description without independent corroboration in our retrieval. The markers are placed in line and in red so that they cannot be missed.

1. Framing Choices, and Where This Inventory Is Weakest

We make three framing choices that a reader should know at the outset. The first concerns separability. A natural way to pose the question is whether frontier AI more closely resembles tobacco, where the harmful element could be engineered away from what users wanted, or asbestos, where the useful property was the harmful property. We find that the answer differs by vector, and that the binary omits the precedent that actually governs the non-separable vectors. Training-data provenance is fully separable, and the courts have already separated it. Open-ended generality and autonomy are not separable: the useful property is the harmful property, which is the asbestos condition. But asbestos, once installed, could not be recalled or rationed, and access to a hosted model can be controlled continuously. The operative precedent for the non-separable vectors is therefore the pharmaceutical regime of restricted distribution, labeled indication and post-market surveillance. Asbestos is what that regime becomes when distribution control fails. Section 6 develops this. The practical consequence is that the checkable question is not whether AI resembles tobacco or asbestos but whether the company describes its distribution control as a risk control whose failure is itself a risk factor.

The second concerns which policy lever matters. For the harm-and-liability risks, the lever that would raise Anthropic's exposure most is the allocation of legal liability for model output. For the firm as a whole that is incomplete, because in 2026 executive discretion moved Anthropic's position faster than any court did. On June 12 the US government, acting through a Commerce Department letter, issued an export-control directive suspending all foreign-national access to Fable 5 and Mythos 5, and Anthropic disabled both models for every customer to comply [S46, S49]. The controls were lifted on conditions on July 1 [S48]. The Department of War's supply-chain-risk designation was held unlawful in the Northern District of California on August 28 [S59] and upheld 2 to 1 by the D.C. Circuit on September 25, the day before our cutoff [S60]. Neither action priced anything. Each withdrew permission. Two levers therefore need watching: liability allocation, which is slow and contested in courts and legislatures, and executive discretion, which is fast and binary.

The third is the scope of the analogy set. The nine industries were chosen to surface harm, reliance and control-regime risks, and they do. They under-generate the commercial risks that will probably dominate the prospectus's risk-factor section by volume: price competition among frontier laboratories, commoditization by open-weight models, rapid obsolescence of each model generation, dependence on key personnel, and the security of model weights and source code. Only telecoms (price collapse after overbuild) and semiconductors (customer concentration) reach into that territory. We do not add a tenth industry. We record the gap so that the scoring rubric can classify these misses as design-induced, and not confuse them with genuine information asymmetry or our own analytical blind spots.

The inventory is weakest in three places. First, every financial figure rests on press-reported run-rates and projections, on analyst estimates of revenue mix, and on a single aggregation of compute commitments by The Information that we could only read secondhand [S31]. No audited statement exists. Second, our assessment of governance rests mainly on Anthropic's own descriptions of its Responsible Scaling Policy, its Long-Term Benefit Trust and its transparency practices [S65, S69, S70]. That is the category of material most heavily represented in the training of the model that drafted this inventory. Our counterweights are one independent rating [S68] and one independent analysis [S71], and the relevant claims are marked accordingly. Third, the harm-from-use risks are populated almost entirely by litigation against OpenAI and Character.AI. We retrieved no wrongful-death or product-liability suit against Anthropic [S98]. Those risks therefore rest on transfer from comparators, not on observation of Anthropic.

Three observable conditions would make the inventory wrong. If courts settle that model output is protected speech or third-party information rather than a product, the harm cluster falls by roughly an order of magnitude. The pleading-stage rulings point the other way: the court in *Garcia* declined to hold chatbot output protected speech [S109], the motion to dismiss in *Lyons v. OpenAI* was denied [S107], and *Moody's* reads Section 230 as unlikely to reach generated content [S114]. None of these is a merits holding. If revenue keeps compounding at the 2026 pace through 2028, the telecom-glut reading is wrong and the capacity risk is shortage rather than surplus. The 2026 outage record [S35] and two consumer class actions over usage limits and degraded service [S95, S96] are evidence of shortage now. If the supply-chain-risk designation is read as narrowly as Anthropic reads it [S58], rather than as broadly as the Secretary's public statement framed it [S57], the procurement risk is about two orders of magnitude below the tail we assign it.

We deliberately left the SASB standards unopened while building the inventory, so that the catalog would not shape it, and mapped the finished inventory against them afterwards; that mapping is reported in a companion working paper. We also deliberately give no estimates of financial impact here. Those are better made with the audited figures the prospectus will contain.

2. The Public Record at the Cutoff

Anthropic, PBC announced on June 1, 2026 that it had confidentially submitted a draft registration statement to the SEC [S01]. As of September 26 no public S-1 appeared on EDGAR [S02]. Press reporting names Morgan Stanley and Goldman Sachs as lead underwriters [S03] and Nasdaq as the venue [S04]. It places the offering in November, so that third-quarter results can be shown, at a valuation near \$2 trillion with proceeds of up to \$100 billion [S05, S06]. The New York Times report of September 18 says Anthropic plans about five gigawatts of compute by the end of 2026 and roughly double that by the end of 2027 [S06]. The founders are reported to be seeking a special share class

carrying 50.1 percent of the vote on most matters. Under that proposal the Long-Term Benefit Trust would continue to select a majority of the board [S07]. The private-market valuation rose from \$183 billion in September 2025 [S08] to \$380 billion in February 2026 [S09] and \$965 billion in May 2026 [S10]. Amazon's second-quarter 10-Q, as summarized, carries its Anthropic interests at about \$190 billion [S19]. Alphabet's non-marketable equity bucket, mainly Anthropic, stood at \$124.3 billion [S20].

The compute position consists of overlapping commitments with parties who are also shareholders. They comprise up to one million Google TPUs [S22] and further multi-gigawatt Google and Broadcom capacity from 2027 [S25]. Anthropic has committed more than \$100 billion to Amazon Web Services over ten years [S26] and \$30 billion to Microsoft Azure alongside Nvidia and Microsoft equity [S24]. The rest is \$50 billion of data centers with Fluidstack [S23], up to two gigawatts of AMD accelerators with an AMD equity investment [S28], a twenty-year lease with TeraWulf [S29], \$45 billion with Nscale [S30], and \$15 billion a year to SpaceX terminable on ninety days' notice [S27]. Chip-lease financing of more than \$71 billion sits in special-purpose vehicles backed in part by Broadcom [S33], and a \$35 billion loan against lease payments carries a Google backstop [S32]. The legal position is summarized in Sections 3 and 4. The final \$450 million installment of the \$1.5 billion Bartz settlement remains outstanding [S87]. Three suits by music publishers seek statutory damages across tens of thousands of works [S89, S90, S92]. The Department of War designation stands after the D.C. Circuit ruling [S60].

3. Cluster I: Harm and Liability Arising from Use

Tobacco

The structural feature doing the work is internal knowledge that ran ahead of public acknowledgment, held by a firm selling under warnings that shifted risk to the user through assumption of risk. The tobacco industry's liability was eventually resolved by settlement with the states, at about \$206 billion over the first twenty-five years of the Master Settlement Agreement [S115]. It was also resolved by a RICO judgment finding a scheme to defraud the public, affirmed on appeal in 2009 [S116]. Later came a harm-reduction regime: a regulator authorized a modified-risk claim once the harmful vector, combustion, was engineered away from the want, nicotine [S117].

For Anthropic the tobacco pattern surfaces first in training-data acquisition, where liability attaches less to the product than to what executives knew and when. The Bartz record established that Anthropic downloaded and retained pirated libraries, and that this was not fair use even though training on lawfully acquired books was [S85]. The music publishers' second suit names Dario Amodei and Benjamin Mann personally and alleges that the torrenting was concealed during discovery in the earlier case [S90]. The Sony and Warner suit extends the same allegations to tens of thousands of further works [S92]. The pattern surfaces second in warnings and assumption of risk. Anthropic's consumer terms cap liability at the greater of six months' fees or \$100 and require users to be eighteen or older. They contain no arbitration clause or class-action waiver, so consumer claims reach court [S100]. The forty-two attorneys general who wrote to Anthropic and twelve other companies in December 2025 demanded safety testing, recall procedures and warnings [S112]; that is the MSA's vocabulary arriving before any litigation. The checkable questions are four. Does the prospectus identify officers implicated in the acquisition conduct at issue, and does it disclose indemnification or advancement of their defense costs? Does it describe user warnings and age assurance? Does it rely on

assumption-of-risk terms? Does the company itself claim to have separated a harm vector from the product's function?

Asbestos

The structural feature is latent injury with a long and indeterminate tail. Under the loss-contingency standard, liability that is probable but not reasonably estimable is disclosed rather than accrued [S121]. When volume overwhelmed estimation, it overwhelmed the balance sheet. Johns-Manville filed under Chapter 11 in 1982 [S120], and about sixty trusts with roughly \$37 billion in assets followed [S118]. The largest trusts paid a median of about a quarter of a claim's stated value [S119], and the insurers withdrew **[unsourced recall]**.

Anthropic has an asbestos-shaped tail in its copyright exposure. The mechanism is not latency of injury but reuse of evidence across plaintiff cohorts. The Bartz release covers only past acquisition conduct and excludes claims based on outputs [S85]. About 350 rightsholders opted out and retain their own claims [S88]. The music publishers' second suit relies expressly on evidence surfaced in Bartz [S90, S91]. Each resolved class seeds the next, which is how asbestos liability propagated. The insurance side of the analogy has already appeared: reporting in October 2025 said insurers were balking at the laboratories' exposures and that Anthropic was partly using its own funds for the Bartz settlement [S101]. On accrual versus disclosure, Bartz is a known quantity. It totals \$1.5 billion, of which \$1.05 billion was paid by August 19, 2026. The final \$450 million plus interest falls due on September 25, 2027 or thirty days after a "qualifying funding event", whichever comes first [S87]. Whether an IPO is a qualifying funding event is unresolved in our sources.

The music suits are where the accrual question is live. At the statutory maximum of \$150,000 per work across more than 20,000 works, the second Concord suit alone carries a ceiling near \$3 billion [S90, S91]. At the Bartz rate of about \$3,000 per work [S86], the same count yields roughly \$64 million **[our estimate: 21,231 works × \$3,000]**. A fifty-fold spread between defensible bounds is exactly the condition under which a preparer concludes that a loss is reasonably possible but not estimable. The checkable questions follow. Does the prospectus accrue anything for the music litigation, or only disclose? Does it state whether the offering accelerates the final Bartz installment? Does it disclose AI-specific liability insurance limits, a captive, or self-insurance?

Pharmaceuticals

The structural feature is that benefit and harm are the same pharmacological property, governed by a sequence of controls. The sequence runs from pre-approval demonstration of safety, through a labeled indication, to mandatory post-market adverse-event reporting on a statutory clock of fifteen calendar days for serious and unexpected events [S122]. Liability attaches to promotion outside the label, as in GlaxoSmithKline's \$3 billion resolution [S123], and to failure to act on a post-market signal, as in Merck's \$4.85 billion Vioxx settlement [S124]. Delay in development is a priced act: firms routinely estimate what a clinical hold costs **[unsourced recall]**.

Anthropic already operates the voluntary skeleton of this regime. Pre-deployment demonstration exists as system cards and as Risk Reports that the current policy schedules every three to six months [S70] **[company statement only]**. A statutory layer now sits alongside it: California's SB 53 requires critical safety incidents to be reported to the Office of Emergency Services [S77], and New York's RAISE Act imposes a seventy-two-hour reporting clock [S78]. Post-market surveillance appears in Anthropic's published misuse reports. These include the August 2025 account of "vibe hacking" extortion against at least seventeen organizations [S102] and the November 2025 account of a state-

sponsored espionage campaign in which the model performed 80 to 90 percent of the work [S103] **[company statement only as to attribution]**. The July 2026 export-control resolution converted part of that practice into an obligation to report malicious activity to authorities [S48].

The off-label analogy is exact in one case. The Commerce directive described its concern as a jailbreak that consisted of asking the model to read a specific codebase and fix software flaws [S46]. That is the labeled indication of Mythos Preview under Project Glasswing, which was offered to defenders and withheld from general availability [S75]. Restricted distribution is visible in three further places: the Annex A list of entities permitted to regain Mythos 5 access [S47], the shelving of an internal model, "Model 2" [S72], and the gating of Glasswing itself. In pharmaceutical terms this is a restricted-distribution program **[unsourced recall: the FDA's REMS mechanism is the reference we have in mind]**. Deliberate slowing is visible in the August 2026 Risk Report's disclosure that "Model 2" will not be released [S72]. The checkable questions are three. Does the prospectus describe a post-deployment incident system with defined reporting clocks and disclose counts of reports filed under SB 53 and RAISE? Does it quantify revenue forgone from withheld or gated models, or the cost of a deployment pause? Does it treat use outside a model's intended deployment as a risk factor distinct from generic misuse?

Social media

The structural feature is intermediary immunity that eroded at the design-defect boundary. Section 230 protects a service from being treated as the publisher of another party's information [S125]. *Lemmon v. Snap* held that a negligent-design claim does not treat the service as publisher and so falls outside the immunity [S126]. From there, design-for-engagement became a product-defect theory. A Los Angeles jury found Meta and YouTube negligent in platform design in March 2026 [S128]. Meta then settled with the state attorneys general for \$12.1 billion guaranteed and up to \$17.1 billion, with mandated age verification and time limits for minors [S129]. Internal research surfaced in discovery and testimony defined what the company knew [S127].

For Anthropic the immunity is probably unavailable from the start, because the model generates the content rather than hosting another party's [S114]. California's SB 243 imposes disclosure and self-harm protocols on companion chatbots, with a private right of action [S113]. Anthropic received both attorney-general letters [S111, S112]. It was not among the seven companies sent FTC 6(b) orders on companion chatbots [S110], which is consistent with its eighteen-plus consumer policy [S100] and an estimated consumer subscription share of 10 to 15 percent of revenue [S12]. We have no public information on whether any Anthropic consumer product optimizes engagement metrics, and we say so rather than infer. The internal-research lesson applies with unusual force, because Anthropic publishes extensively about its own models' behavior, including usage research that measures how autonomously its products are used [S147] **[unsourced recall as to the breadth of its model-behavior publications]**. What a company publishes about its product becomes the benchmark of what it knew. The checkable questions are four. Does the prospectus claim any intermediary immunity or disclaim it? Does it describe age assurance, including for minors reached through third-party applications built on the API? Does it identify published internal research as a litigation risk? Does it disclose any engagement-optimization objective for consumer products?

4. Cluster II: Reliance, Intermediation and Contractual Allocation

Credit rating agencies

The structural feature is a payment triangle. The issuer pays, the investor relies, and the SEC has long acknowledged the conflict of interest the model creates [S131]. The agencies' opinion defense eroded where ratings went to a select group rather than the public [S134]. Regulation embedded reliance on ratings and then, after the failure, removed it through section 939A and stripped the expert-liability exemption through section 939G [S131]. The settlements rested on internal evidence that business relationships had shaped model decisions: S&P paid \$1.375 billion [S132], and Moody's paid about \$864 million with CEO certifications [S133].

Anthropic sits inside two triangles. The first is the deployment triangle. An enterprise pays, and that enterprise's customers, patients, borrowers and counterparties rely on outputs without any contract with Anthropic. The commercial terms cap Anthropic's liability at twelve months' fees and exclude consequential damages, while its intellectual-property indemnity to paying customers sits outside the cap [S99]. On the opinion defense, the rating-agency cases suggest that the protection is weakest where output is distributed privately. Enterprise API output is private by construction.

The second is the self-assessment triangle. Anthropic determines whether its own models cross the capability thresholds in its Responsible Scaling Policy, a rater rating itself, which is structurally weaker than issuer-pays. The Trust can mandate external review but had not used that power by August 2026 [S72]. An independent analysis finds that the version 3 policy triggers external review only narrowly [S71], and an independent rating scores Anthropic's risk management at 50 percent [S68]. The regulatory reliance of the pre-2008 rating agencies has an analog here. SB 53, the RAISE Act and the EU Code of Practice build on developer-authored frameworks [S77, S78, S80], and the July 2026 Commerce conditions rest on Anthropic's own detection of security risks [S48]. The post-failure precedent says reliance is withdrawn and shields are stripped once internal evidence shows commercial considerations shaped the model. The sentence most likely to be read as such evidence is already public. It is Jared Kaplan's explanation of the version 3 policy: it did not make sense to make unilateral commitments "if competitors are blazing ahead" [S146]. The checkable questions are three. Does the prospectus describe any external, independently paid assessment of model risk, and name who pays for it? Does it disclose that regulatory regimes rely on its self-assessment, and treat the withdrawal of that reliance as a risk? Does it characterize model output as opinion or information, with a disclaimer of third-party reliance?

ERP and enterprise software

The structural feature is that value depends on the customer's deployment, so failure of the implementation becomes a contest over pre-sale representations. Waste Management sued SAP alleging demonstrations run on fake software environments, seeking more than \$100 million [S135]. Revlon disclosed a material weakness from an SAP go-live, and shareholder class actions followed [S136]. Lidl abandoned an SAP program after about €500 million [S137]. The industry's mature response is contractual. Liability is capped at annual fees, consequential damages are excluded, with specified carve-outs [S138].

Anthropic's commercial terms replicate that architecture almost clause for clause [S99]. The analogy therefore does its work at the two points where the architecture is untested. First, the agentic product acts. Claude Code executes in customer environments, so the ERP failure mode of a mis-configured system becomes an autonomous action with consequences, and fee-based caps have not yet been tested against an agent-caused loss. Second, the representation risk is the ERP pattern on a larger scale. A

widely cited study found that 95 percent of organizations reported no profit-and-loss return from generative-AI spending [S148], the condition under which ERP customers turned to litigation. The ERP precedent also points to customer-side disclosure. For bank customers the natural channel for a Revlon-type failure would have been supervisory model validation. However, the US banking agencies' revised model-risk guidance of April 2026 explicitly excludes generative and agentic AI models from its scope and leaves banks to apply their general risk-management practices to them [S150, S151]. The customer-side check on a vendor model is therefore now governed by no specific supervisory standard. The checkable questions are three. Are the limitation-of-liability terms and the uncapped indemnity disclosed, with any quantification of claims made? Does the prospectus disclose customer disputes or credits arising from agentic actions? Does it describe how customers' model-validation requirements, including notice of model version changes, are met?

Telecommunications

Three structural features do the work. The first is capital intensity followed by a capacity glut. Communications-equipment investment roughly doubled between 1996 and 2000, and telecom firms then lost about \$700 billion of market value [S139]. At the trough, 97 percent of fibre sat unused at any given moment [S140]. WorldCom capitalized about \$3.8 billion of operating costs [S141], and equipment makers financed their own customers ***[unsourced recall: the Lucent and Nortel vendor-financing precedent]***. The second is the distinction between conduit and publisher. A common carrier carries all lawful traffic, whereas a publisher exercises editorial choice. The third is the path by which essential infrastructure acquires price and access regulation, contested in the United States through the 2025 Sixth Circuit decision on broadband [S142].

For Anthropic the capacity analogy is the most consequential in the inventory. The company plans about five gigawatts by the end of 2026 and about ten by the end of 2027 [S06], against commitments aggregated at about \$517 billion [S31]. The duration structure is mismatched. It combines a twenty-year lease [S29], chip-lease SPVs [S33], and a \$15 billion annual contract terminable on ninety days [S27], all serving revenue with a history measured in quarters. Allianz estimates that AI investment is growing about 46 percent faster than revenue, against 32 percent in the 2001 telecom excess [S42]. Bain estimates an \$800 billion annual revenue shortfall by 2030 [S44]. Goldman Sachs identifies chip useful life as the single most influential variable in the build-out's economics [S43]. The difference from telecom lies in where the debt sits. The telecom builders borrowed on their own balance sheets. Here much of the debt sits with lessors and SPVs, guaranteed by Google and Broadcom [S32, S33], and Anthropic's exposure takes the form of take-or-pay obligations. The analogy also cuts the other way today, because the current problem is shortage [S35, S95, S96].

The conduit analogy clarifies the dispute with the Department of War. The Pentagon's demand that Claude be available for "all lawful use" [S60] is a common-carriage demand. Anthropic's refusal to permit mass domestic surveillance and fully autonomous weapons [S56] is an editorial choice of the kind a publisher makes. The access-regulation path is already visible. Commerce decided which entities could use Mythos 5 [S47], and the July conditions included approval to serve critical-infrastructure operators [S48], an allocation of access by the state. Energy supplies the rate-base leg: a record PJM capacity price driven by data-center demand [S38], Anthropic's pledge to cover grid upgrades and consumer price effects [S37], and criticism of gas turbines at a leased site [S39]. The checkable questions are four. Does the prospectus disclose purchase and lease obligations by year, with cancellation terms and guarantor identity? What share of capacity sits in structures off Anthropic's balance sheet? Is there a sensitivity of fixed compute obligations to revenue shortfall? Does the

company treat its acceptable-use restrictions as a possible basis for non-discrimination or access regulation?

5. Cluster III: Systemic Exposure and Control Regimes

Banking

The structural feature is a supervised regime for model risk. The 2011 supervisory guidance defined model risk as adverse consequences from decisions based on incorrect or misused model outputs, and required independent validation and "effective challenge" [S143]. It was superseded on April 17, 2026 by revised guidance. That guidance retains effective challenge in a less prescriptive form and explicitly excludes generative and agentic AI models from its scope, on the ground that they are novel and rapidly evolving [S150, S151]. The supervisory regime that anchors this analogy has thus been drawn, at the moment of Anthropic's offering, so as to leave out exactly the technology in question. Large banks must maintain a risk committee chaired by an independent director and a chief risk officer who reports directly to that committee as well as to the chief executive [S144]. Stress testing discloses tails that cannot be quantified precisely, and the Federal Reserve is now moving to publish its loss models [S145].

Anthropic's Responsible Scaling Policy is a model-risk framework with no external supervisor. The independent rating finds a designated Responsible Scaling Officer but no central risk function, a board risk committee with substantial gaps, and limited internal audit [S68]. Anthropic has said that changes to the policy require board approval [S149] *[company statement only]*. The Trust-selected directors became a board majority in April 2026 [S66]. In August a trustee left the Trust to become the company's chief global affairs officer [S67], a move from overseer to management that banking governance would treat as a break in independence. The proposed founder voting control [S07] adds a controlled-company layer above the board. The nearest thing to a stress disclosure is the August 2026 Risk Report, which raised the assessed misalignment risk to "low" [S72]. We found no public stress test of fixed compute obligations against revenue. The checkable questions are four. Is there a chief risk officer, or an equivalent, reporting directly to a board risk committee chaired by an independent director? Is the Responsible Scaling Officer independent of the chief executive and the product line? Does the prospectus present a stress scenario for revenue against fixed obligations? Are Responsible Scaling Policy commitments described as binding, and does it say who may amend them?

Semiconductors

The structural features are disclosed concentration, export controls with direct income-statement consequences, and chokepoint disclosure. Nvidia reports its largest direct customers at 22 and 14 percent of revenue and describes itself as effectively foreclosed from China's data-center market [S55]. It took a \$4.5 billion charge when H2O exports were restricted [S54]. TSMC's ten largest customers account for 78 percent of its revenue, and it discloses conditions in Taiwan as a risk [S53].

For Anthropic the analogy exposes concentration in which the suppliers are also investors. Amazon's investment facility is tied to compute-delivery milestones [S19], and Anthropic's commitment of more than \$100 billion flows back to AWS [S26]. Google is investing up to \$40 billion, holding about 14 percent capped at 15 [S21], while supplying TPUs [S22] and guaranteeing lease financing [S32]. Microsoft and Nvidia invested alongside a \$30 billion Azure commitment [S24], AMD's equity investment accompanies its accelerator supply [S28], and Broadcom backs SPV tranches [S33]. The investors book gains from marks on Anthropic [S19, S20] while Anthropic's spending appears as their

cloud revenue. This is circularity of the kind that analysts have begun to flag across the sector [S41]. Nvidia's chief executive has said its investments in the leading laboratories are likely to be its last [S40]. On export controls, Anthropic has experienced a model-level control directly [S46]. It advocates chip-level controls [S50, S51], restricts sales to Chinese-controlled entities [S45], and sits downstream of the Taiwan chokepoint [S53]. On the customer side, Cursor and GitHub Copilot were reported at nearly a quarter of revenue in 2025 [S13]. The checkable questions are four. Does the prospectus name any customer above 10 percent of revenue, and any supplier above 10 percent of cost of revenue? Are compute contracts with shareholders disclosed as related-party transactions, with any milestone link between investment tranches and compute purchases? Is revenue lost during the June 2026 suspension quantified? Is Taiwan named as a chokepoint?

6. Separability

Tobacco harm reduction works because the want is nicotine and the harm is combustion, so delivery can be re-engineered once the vector is identified. The FDA's modified-risk order for a heated product is the regulatory form of that separation [S117]. Asbestos had no harm-reduced version, because the useful property was the harmful property. We test the four candidate vectors that matter most (training-data provenance, the anthropomorphic interface, open-ended generality and autonomy) and reach different answers for each.

Training-data provenance is separable, and the separation has been demonstrated. The Bartz rulings held that training on lawfully acquired books was transformative fair use and that acquiring and keeping pirated copies was not [S85]. The harm was the method of acquisition, not the training. Re-engineering means purchase and licensing, which is the tobacco case exactly. On the output side, lyric reproduction is partly separable through guardrails, which were stipulated in the first Concord case [S89]. This vector is governed by tobacco and is already in the harm-reduction phase. What remains is legacy liability for past conduct, and in tobacco that is where the money was.

The anthropomorphic interface is separable in principle and contested in practice. Disclosure of AI status, crisis protocols and age gates, the substance of SB 243 [S113], can be added at modest cost without removing function. Anthropic's eighteen-plus consumer policy [S100] is itself a separation. But part of what users want from a conversational agent is its conversational character. The vector is therefore only partly separable, and the unseparated remainder is exactly where the social-media design-defect theory would attach.

Open-ended generality is not separable. The Commerce directive is the proof: the capability it treated as a harm, reading a codebase and fixing its flaws [S46], is the capability Glasswing sells to defenders [S75]. This is the asbestos condition. It differs from asbestos in one decisive respect. Asbestos, once installed, could not be recalled, whereas access to a hosted model can be rationed account by account, list by list and version by version. Anthropic's actual responses to this vector are pharmaceutical: restricted release [S75], an access list [S47], a withheld model [S72], and surveillance and reporting duties [S48]. The state has adopted the same frame.

Autonomy is not separable either. In the espionage campaign Anthropic disclosed, the model performed 80 to 90 percent of the work [S103]. Anthropic's own usage research measures its coding product as used more autonomously than chat [S147]. The value is action without a human, and so is the harm. Mitigation through monitoring and human checkpoints trades directly against the product's function.

Our position follows. The governing precedent is decided vector by vector, and Anthropic's portfolio is split. The liability that has crystallized, copyright, sits on the one fully separable vector. It is tobacco-shaped: large, legacy, finite and re-engineerable. The exposure that has not crystallized sits on the two non-separable vectors. It is asbestos-shaped in its underlying property but pharmaceutical-shaped in governance, and it stays pharmaceutical only while distribution control holds. Three developments would convert it to asbestos: exfiltrated weights, open-weight equivalents at the frontier, or distribution through channels that the developer cannot monitor. That makes one question in the prospectus more diagnostic than any other. Does Anthropic describe its gating (restricted access, customer screening, usage monitoring) as a control whose failure is a risk factor, and does it say what happens to its liability position if equivalent capability becomes available without gating?

7. The Consolidated Inventory

The table deduplicates the risks surfaced by the nine industries into fifteen risks, the last of which is a question of disclosure rather than exposure. Where several industries reach the same risk by different routes, we record the routes. We treat convergence as evidence, because independent structural analogies arriving at the same exposure make it less likely that any one analogy is doing the work alone.

#	Risk	Generating industries	Routes of convergence
1	Compute commitment mismatch (glut later, shortage now)	Telecoms, semiconductors, banking	Capacity overbuild and duration mismatch; counterparty obligations; absence of stress disclosure
2	Energy, grid and host-community exposure	Telecoms	Infrastructure acquiring rate-base and access regulation
3	Supplier–investor concentration and circular financing	Semiconductors, telecoms, credit rating agencies	Supplier concentration; vendor financing; conflicted payment relationships
4	Customer and channel concentration	Semiconductors, telecoms	Disclosed customer concentration; price collapse on loss of volume
5	Revenue quality and presentation (run-rate versus GAAP, gross versus net through cloud marketplaces, related-party revenue)	Telecoms, ERP, semiconductors	Capitalization and presentation abuse after overbuild; customer-side disclosure; circular revenue
6	Sovereign access discretion (model- and chip-level export controls, Taiwan chokepoint)	Semiconductors, telecoms, pharmaceuticals	Export-control charges; state allocation of access; regulator-authorized distribution
7	Procurement designation and political exposure	Telecoms, banking	Common-carriage demand against editorial refusal; supervisory discretion
8	Training-data acquisition liability (legacy piracy, follow-on suits, officer exposure)	Tobacco, asbestos, social media	Internal knowledge of executives; tail propagated by reuse of evidence across cohorts; discovery surfacing internal conduct
9	Output IP and customer indemnity	ERP, credit rating agencies	Contractual allocation of third-party claims; uncapped indemnity inside capped terms
10	Harm from use: psychological harm, minors, dependency	Tobacco, social media, pharmaceuticals	Warning and assumption of risk; design defect outside intermediary immunity; adverse-event duty
11	Third-party misuse of dual-use capability (cyber, bio) and uninsurability of the tail	Pharmaceuticals, asbestos, telecoms, semiconductors	Off-label use; indeterminate tail and insurer exit; conduit versus publisher; export-control framing of capability
12	Agentic deployment and implementation failure; third-party reliance on output	ERP, credit rating agencies, banking	Implementation litigation; payment triangle and opinion defense; customer model validation
13	Regulatory regime proliferation and post-failure discontinuity	Pharmaceuticals, social media, credit rating agencies, banking	Reporting clocks; minors' regimes; reliance withdrawn after failure; supervision

#	Risk	Generating industries	Routes of convergence
14	Deployment restraint as a priceable act (withheld or gated models)	Pharmaceuticals, credit rating agencies	Priced development delay; competitive pressure on self-imposed standards
15	Risk-oversight architecture and control structure (risk function, Trust, founder control)	Banking, credit rating agencies, social media, pharmaceuticals	Independent chief risk officer; self-assessment; controlled company overriding internal research; safety function

Table 1. Consolidated analogical risk inventory

Risks 11, 13 and 15 are each reached by four industries through distinct mechanisms, and risks 1, 3, 6, 8, 10 and 12 by three. We single out risk 8 because its convergence has already been confirmed: the tobacco route (what executives knew), the asbestos route (evidence reused by successive plaintiffs) and the social-media route (discovery) are each visible in the second Concord complaint [S90]. One variable is reached by four industries at once and is better treated as a lever than as a position: the legal characterization of model output. Tobacco asks whether output is a product with a warning. Social media asks whether the developer is the publisher of its own output. The rating agencies ask whether output is protected opinion. Telecoms asks whether the developer is a conduit or an editor. The answer to that question determines the vertical position of risks 10, 11 and 12 together, which is why we treat it as a lever rather than a separate risk.

The residual risks the analogies did not generate are price competition among frontier laboratories, commoditization by open-weight models, obsolescence of each model generation, dependence on key personnel, and the security of weights and code. Of these, the only public incident we retrieved is the March 2026 leak of Claude Code source [S97]. We expect these to account for much of the prospectus's risk-factor section, for the reason given in Section 1.

We did not open the SASB standards for this pass. Working from recall of the taxonomy's structure **[unsourced recall]**, we expect the eventual mapping to touch roughly twelve to fifteen of the seventy-seven industries. Software and IT Services and Internet Media and Services carry most of the inventory. Semiconductors, Telecommunication Services, Electric Utilities, Biotechnology and Pharmaceuticals, the banking industries (for model risk), Media and Entertainment (for intellectual property) and Real Estate (for leased data-center capacity) carry the remainder. We expect a substantial minority of risks to have no home at all. The likeliest candidates are separability, sovereign access discretion, deployment restraint and the self-assessment triangle. That count will be a finding about the standards rather than about Anthropic.

8. Policy Advocacy and Its Consistency

For risks that sit largely outside the company's control, the natural question is what public policy the company advocates to address them, and whether its lobbying is consistent with that advocacy. We apply that test to compute commitments and energy (1, 2), supplier and customer concentration (3, 4), export controls (6), the procurement designation (7) and dual-use misuse (11), and to one position that Anthropic has taken about its own commitments.

On dual-use misuse, Anthropic advocates transparency legislation. It endorsed SB 53 [S76], supported the RAISE Act [S78], opposed a ten-year federal moratorium on state laws as "too blunt" [S79], and funded Public First Action with \$40 million on a platform of frontier transparency, a federal framework that does not preempt stronger state laws, chip export controls and rules targeting biological and cyber risks [S83, S84]. Its federal lobbying reached \$3.1 million in 2025 and more than \$3.5 million in the first half of 2026, on export controls, procurement, supply-chain designation and

its acceptable-use policy among other issues [S81, S82]. Lobbying and advocacy are consistent on transparency. The gap is liability allocation. We retrieved no public position from Anthropic on who should bear liability for harm caused through its models, whether safe harbor, negligence or strict liability. Yet that is the variable that would most change its exposure. The test therefore returns an absence, and the absence is the finding.

On compute commitments and energy, Anthropic advocates rapid build-out, citing a national need of at least fifty gigawatts by 2028 and calling for permitting reform [S36], and it has pledged to absorb grid-upgrade costs [S37]. Advocacy and commitments point the same way. On supplier-investor concentration and circular financing, we found no advocacy at all; control sits with Amazon, Google, Nvidia and Broadcom, and no public policy position addresses the arrangement. On customer concentration the test does not apply, because the risk is commercial and has no policy counterpart. Not every risk outside a company's control has a public-policy hand-off.

On export controls, Anthropic advocates strict controls on AI chips. It called H200 sales to China "a big mistake" [S50] and argued for a tighter diffusion rule over Nvidia's objection [S51]. At the same time it contested the model-level directive applied to itself, arguing that the standard would halt deployments across the industry [S46]. The positions are not inconsistent in logic, since controls on hardware and controls on hosted services differ. But both coincide with Anthropic's competitive position, a familiar pattern in which the policy a company advocates favors its own portfolio. The charge of regulatory capture by the White House AI adviser David Sacks [S63] is the adversarial reading of the same facts. On the procurement designation, Anthropic's advocacy (its two use restrictions), its litigation and its lobbying on procurement and supply-chain risk [S81] are consistent with one another, and it has accepted the cost of that consistency.

The last case is a position Anthropic has taken about its own commitments. Version 3 of the Responsible Scaling Policy removed the pause commitment and downgraded the SL-4 security standard from a unilateral commitment to an industry recommendation [S71], separating what Anthropic will do alone from what it recommends for the industry because stronger mitigations might prove impossible to implement without collective action [S70]. That moves those issues from the company's own control to public policy. The test then asks whether Anthropic advocates binding versions of what it moved out of unilateral commitment. The Public First Action platform covers transparency and targeted biological and cyber rules [S83]. We found no advocacy for a mandated pause mechanism or a mandated security level. The consistency is partial, and the question is well posed for the prospectus.

9. The Prospectus Question Set

Each question below is written so that it can be scored against the filed prospectus. A scoring protocol, with hit, partial and miss defined, is provided in a companion working paper for readers who wish to make the comparison systematically. The governing premise stands: a risk-factor section is a legal artifact and not a materiality assessment, so coverage in the prospectus is not the ground truth of materiality.

Q	Question	Form	Risk	Generated by
1	Does the prospectus accrue any amount for the music-publisher litigation, or disclose it as reasonably possible but not estimable?	Accrue / disclose / silent	8	Asbestos

Q	Question	Form	Risk	Generated by
2	Does it state whether the offering is a qualifying funding event that accelerates the final Bartz installment?	Yes / no	8	Asbestos
3	Does it disclose indemnification or advancement of defense costs for officers named personally in litigation?	Yes / no	8	Tobacco
4	Does it disclose AI-specific liability insurance limits, a captive, or self-insurance?	Quantified / described / silent	8, 10, 11	Asbestos
5	Does it describe a post-deployment incident or adverse-event system with reporting clocks, and give counts of reports under SB 53 or RAISE?	Counts / described / silent	10, 11, 13	Pharmaceuticals
6	Does it quantify revenue forgone from withheld or gated models, or the cost of a deployment pause?	Quantified / described / silent	14	Pharmaceuticals
7	Does it describe access gating as a risk control whose failure (weight exfiltration, open-weight equivalents) is itself a risk factor?	Yes / partial / no	11	Pharmaceuticals, asbestos (separability)
8	Does it claim, disclaim or stay silent on intermediary immunity for model output?	Claim / disclaim / silent	10	Social media
9	Does it describe age assurance, including minors reached through third-party applications on the API?	Yes / partial / no	10	Social media
10	Does it identify its own published research on model behavior as a potential source of litigation risk?	Yes / no	10, 15	Social media, tobacco
11	Does it disclose its limitation-of-liability terms and the uncapped IP indemnity, with any quantification of claims?	Quantified / described / silent	9, 12	ERP
12	Does it disclose disputes or credits arising from agentic actions in customer environments?	Yes / no	12	ERP
13	Does it describe any external, independently paid assessment of model risk, and who pays?	Yes / partial / no	15	Credit rating agencies
14	Does it disclose that regulatory regimes rely on its self-assessment, and the risk that such reliance is withdrawn?	Yes / no	13, 15	Credit rating agencies
15	Is there a chief risk officer or equivalent reporting directly to a board risk committee chaired by an independent director?	Yes / partial / no	15	Banking
16	Are Responsible Scaling Policy commitments described as binding, and who may amend them?	Binding / discretionary / silent	14, 15	Banking, pharmaceuticals
17	Does it present a stress scenario of revenue against fixed compute obligations?	Quantified / qualitative / silent	1	Banking, telecoms
18	Does it disclose purchase and lease obligations by year, with cancellation terms, guarantor identity and off-balance-sheet share?	Full / partial / silent	1	Telecoms
19	Are compute contracts with shareholders disclosed as related-party transactions, with any milestone link to investment tranches?	Yes / partial / no	3	Semiconductors, telecoms
20	Does it name customers above 10 percent of revenue and suppliers above 10 percent of cost of revenue?	Named / anonymized / silent	3, 4	Semiconductors
21	Does it quantify the revenue effect of the June 2026 export-control suspension?	Quantified / described / silent	6	Semiconductors
22	Does it name Taiwan or a single foundry as a chokepoint?	Yes / no	6	Semiconductors
23	How does it describe the scope of the supply-chain-risk designation after the D.C. Circuit ruling, and does it state whether any partner has been asked to cease commercial activity?	Broad / narrow / silent	7	Telecoms
24	Does it state a public-policy position on liability allocation for harm caused through its models?	Yes / no	10, 11, 12	Accountability companion
25	Does it quantify the electricity-price and grid-upgrade pledge as a commitment?	Quantified / described / silent	2	Telecoms
26	Does it reconcile run-rate to GAAP revenue, and state gross or net recognition for sales through cloud marketplaces?	Both / one / neither	5	Telecoms, ERP

Table 2. Checkable questions for the prospectus

Appendix A. Source Register

All entries were retrieved on September 26, 2026 from sources published on or before that date. Type distinguishes company statements, regulatory filings (primary or summarized), court documents, government sources, press reporting, and analyst, secondary-market or independent estimates. Where the register relies on a secondary summary of a primary document, the type says so. Conflicts noted during retrieval: the Bartz final-approval date is reported as July 20 and July 22, 2026; the Department of War designation date is reported as late February and March 4, 2026; the court issuing the September 19 Reddit ruling is ambiguous; revenue and year-end projections vary across outlets (above \$100 billion versus above \$110 billion). The Amazon and Alphabet percentage stakes rest on a single report and were not used.

ID	Claim	Source	URL	Published	Type
So1	Anthropic, PBC confidentially submitted a draft Form S-1; share count and price not set	Anthropic, "Anthropic confidentially submits draft S-1"	https://www.anthropic.com/news/confidential-draft-s1-sec	Jun 1, 2026	Company statement
So2	No public Anthropic S-1 on EDGAR as of Sep 26, 2026	Forkast, "The Anthropic S-1 that was expected after Labor Day still hasn't arrived"	https://forkast.news/the-anthropic-s-1-that-was-expected-after-labor-day-still-hasnt-arrived/	Sep 26, 2026	Press reporting
So3	Morgan Stanley and Goldman Sachs to co-lead IPO; JPMorgan participating	Bloomberg via Yahoo Finance	https://finance.yahoo.com/markets/stocks/articles/anthropic-said-pick-morgan-stanley-175802678.html	Jun 3, 2026	Press reporting
So4	Anthropic chose Nasdaq as listing venue	Bloomberg via Yahoo Finance	https://finance.yahoo.com/markets/stocks/articles/anthropic-said-choose-nasdaq-much-180359465.html	Sep 13, 2026	Press reporting
So5	IPO moved to November to show Q3 results; ~\$2T valuation; raise up to \$100B (WSJ)	Investing.com	https://www.investing.com/news/stock-market-news/anthropic-delays-ipo-staging-to-november-amid-ai-fears-wsj-says-4907941	c. 18–Sep 19, 2026	Press reporting
So6	Financial documents public in coming weeks; trading as soon as November; ~\$2T; >\$100B annualized revenue by YE 2026 vs \$65B in July; ~5 GW compute by YE 2026, roughly double by YE 2027 (NYT, 18 Sep)	DT Next (NYT republication); Newsquawk headline	https://www.dtnext.in/news/business/growth-slowdown-dynamics-anthropic-pursues-ipo-amid-ai-safety-debate ; https://www.newsquawk.com/headlines/anthropic-reportedly-plans-to-have-5gw-worth-of-compute-by-the-end-of-2026-reports-nyt	18–Sep 21, 2026	Press reporting
So7	Founders seek special shares with 50.1% of vote on most matters; LTBT continues to select most directors; founder seats 2→3; Aug 2026 secondary valuation ~\$1.5T (The Information)	TechCrunch	https://techcrunch.com/2026/09/25/anthropics-founders-seek-voting-control-ahead-of-ipo/	Sep 25, 2026	Press reporting
So8	Series F: \$13B at \$183B post-money	Anthropic	https://www.anthropic.com/news/anthropic-raises-series-f-at-usd183b-post-money-valuation	Sep 2, 2025	Company statement
So9	Series G: \$30B at \$380B; run-rate \$14B; Claude Code run-rate >\$2.5B	Anthropic	https://www.anthropic.com/news/anthropic-raises-30-billion-series-g-funding-380-billion-post-money-valuation	Feb 12, 2026	Company statement
So10	Series H: \$65B at \$965B; \$15B from hyperscalers; run-rate >\$47B in May 2026	Anthropic	https://www.anthropic.com/news/series-h	May 28, 2026	Company statement

ID	Claim	Source	URL	Published	Type
S11	Annualized revenue >\$65B at end-July 2026	TechCrunch (following Bloomberg)	https://techcrunch.com/2026/08/17/anthropics-annualized-revenue-surges-to-65b/	Aug 17, 2026	Press reporting
S12	Revenue mix estimate: ~70–75% API, 10–15% consumer subscriptions	Sacra	https://sacra.com/c/anthropic/	2026 (undated)	Analyst estimate
S13	Cursor and GitHub Copilot ~\$1.2B, nearly a quarter of a ~\$5B run-rate	VentureBeat	https://venturebeat.com/business/anthropic-revenue-tied-to-two-customers-as-ai-pricing-war-threatens-margins	Aug 8, 2025	Press reporting / analyst estimate
S14	Q2 2026 revenue ~\$10.9B projected with first operating profit; may not remain profitable through year (WSJ)	Yahoo Finance	https://finance.yahoo.com/sectors/technology/articles/anthropic-says-first-profitable-quarter-002121100.html	May 20, 2026	Press reporting of company projection
S15	Q1 2026 revenue \$4.8B (WSJ); Q2 \$559M operating profit projection described as non-GAAP; inference costs 23% above plan (The Information)	Ed Zitron, Where's Your Ed At	https://www.wheresyoured.at/anthropics-profitability-swindle/	2026	Analyst estimate (critical secondary)
S16	Projected gross margin ~94% (2024), ~50% (2025), 77% (2028); 2028 revenue \$70B	TechCrunch (The Information)	https://techcrunch.com/2025/11/04/anthropic-expects-b2b-demand-to-boost-revenue-to-70b-in-2028-report/	Nov 4, 2025	Press reporting of company projections
S17	2025 gross margin projection lowered to 40%	The Information (paywalled; headline and lede only)	https://www.theinformation.com/articles/anthropic-lowers-profit-margin-projection-revenue-skyrockets	early 2026	Press reporting
S18	\$15B pre-IPO revolving credit facility led by Morgan Stanley	Bloomberg via Yahoo Finance	https://finance.yahoo.com/technology/ai/articles/anthropic-nears-finalizing-15-billion-224354345.html	Sep 3, 2026	Press reporting
S19	Amazon Q2 2026 10-Q: \$50.5B fair-value gain; Anthropic interests ~\$190.4B; \$20B financing facility (\$15B available) tied to compute-delivery milestones	StockTitan summary of Amazon 10-Q	https://www.stocktitan.net/sec-filings/AMZN/10-q-amazon-com-inc-quarterly-earnings-report-ob79ef7cd1d2.html	Jul 30, 2026	Regulatory filing (secondary summary)
S20	Alphabet non-marketable equity, mainly Anthropic, \$124.3B at Q2 2026	Motley Fool	https://www.fool.com/investing/2026/09/24/alphabet-values-its-anthropic-stake-at-usd124-billion-a-usd2-trillion-ipo-could-double-it-on-paper/	Sep 24, 2026	Analyst estimate (of filing)
S21	Google to invest up to \$40B; owns ~14%, capped at 15%; Amazon \$16.8B Q1 pre-tax gain	Fortune	https://fortune.com/2026/06/04/amazon-google-billions-anthropic-ipo/	Jun 4, 2026	Press reporting
S22	Up to one million Google TPUs; >1 GW in 2026; tens of billions of dollars	Anthropic	https://anthropic.com/news/expanding-our-use-of-google-cloud-tpus-and-services	Oct 23, 2025	Company statement
S23	\$50B US data-center investment with Fluidstack (Texas, New York)	Anthropic	https://www.anthropic.com/news/anthropic-invests-50-billion-in-american-ai-infrastructure	Nov 12, 2025	Company statement

ID	Claim	Source	URL	Published	Type
S24	\$30B Azure commitment, up to 1 GW; Microsoft up to \$5B and Nvidia up to \$10B investment	Anthropic	https://www.anthropic.com/news/microsoft-nvidia-anthropic-announce-strategic-partnerships	Nov 18, 2025	Company statement
S25	Multiple GW of Google/Broadcom TPU capacity from 2027; run-rate >\$30B; >1,000 customers >\$1M/yr	Anthropic	https://www.anthropic.com/news/google-broadcom-partnership-compute	Apr 6, 2026	Company statement
S26	>\$100B committed to AWS over ten years, up to 5 GW; Amazon \$5B plus up to \$20B	Anthropic	https://www.anthropic.com/news/anthropic-amazon-compute	Apr 20, 2026	Company statement
S27	Anthropic pays SpaceX \$1.25B/month through May 2029; either party may exit on 90 days' notice (SpaceX S-1)	Axios	https://www.axios.com/2026/05/20/anthropic-spacex-compute	May 20, 2026	Regulatory filing (third-party) via press
S28	Up to 2 GW AMD MI450; AMD equity investment up to \$5B	AMD Investor Relations	https://ir.amd.com/news-events/press-releases/detail/1292/amd-and-anthropic-announce-strategic-partnership-to-deploy-up-to-2-gigawatts-of-amd-instinct-mi450-series-gpus	Jul 22, 2026	Company statement (counterparty)
S29	20-year TeraWulf lease, ~401 MW, ~\$19B contracted revenue	TeraWulf Investor Relations	https://investors.terawulf.com/news-events/press-releases/detail/142/terawulf-announces-anthropic-lease-at-justified-data-campus-and-sale-of-majority-interest-in-abernathy-joint-venture-to-fluidstack	Jul 6, 2026	Company statement (counterparty)
S30	\$45B six-year deal with Nscale, ~460 MW	TechCrunch	https://techcrunch.com/2026/08/26/anthropic-continues-compute-gobbling-streak-in-45-billion-deal-with-nscale/	Aug 26, 2026	Press reporting
S31	Compute agreements totalling ~\$517B (about 14.8 GW) signed in the eleven months since October 2025, compiled by The Information from announcements and its own reporting	DatacenterDynamics (reporting The Information); 24/7 Wall St.	https://www.datacenterdynamics.com/en/news/anthropic-signed-517bn-in-compute-agreements-in-past-11-months/ ; https://247wallst.com/investing/2026/09/13/anthropic-locks-down-517-billion-in-compute-ahead-of-ipo-and-its-still-not-enough/	Sep 7, 2026; Sep 13, 2026	Press reporting (secondary aggregation)
S32	\$35B loan against data-center lease payments; Google backstop; Apollo and Blackstone lenders	Business Standard	https://www.business-standard.com/world-news/google-helps-secure-35-billion-funding-for-anthropic-data-centers-126061001276_1.html	Jun 10, 2026	Press reporting
S33	Chip-lease SPVs >\$71B combined, tranches backed by Broadcom; site financing estimates	Yahoo Finance; Epoch AI	https://finance.yahoo.com/technology/ai/articles/anthropic-spvs-stack-71-billion-000514097.html ; https://epochai.substack.com/p/will-financing-bottleneck-ai-compute	Aug 2026	Press reporting / analyst estimate
S34	Talks to lease up to 1 GW from Stream Data Centers; Google may guarantee (The Information)	Yahoo Finance	https://finance.yahoo.com/technology/ai/articles/anthropic-discusses-1-gw-stream-134246854.html	Sep 23, 2026	Press reporting
S35	164 Claude incidents since Jan 2026; "unexpected capacity constraints"	TechTimes	https://www.techtimes.com/articles/323171/20260805/claude-goes-down-again-71b-compute-deal-cannot-prevent-anthropics-164th-outage.htm	Aug 5, 2026	Press reporting
S36	US AI sector needs ≥50 GW by 2028; call for permitting reform	Anthropic, "Build AI in America"	https://anthropic.com/news/build-ai-in-america	Jul 21, 2025	Company statement

ID	Claim	Source	URL	Published	Type
S37	Pledge to pay 100% of grid upgrades and cover consumer electricity price effects	Anthropic, "Covering electricity price increases"	https://anthropic.com/news/covering-electricity-price-increases	Feb 11, 2026	Company statement
S38	PJM 2027/28 capacity auction record \$333.44/MW-day, data-center driven	Utility Dive	https://www.utilitydive.com/news/pjm-interconnection-capacity-auction-data-center/808264/	Dec 18, 2025	Press reporting
S39	Colossus 1 gas turbines criticized by civil-rights groups	DatacenterDynamics	https://www.datacenterdynamics.com/en/news/anthropic-to-use-all-of-spacex-xais-colossus-1-data-center-compute/	May 6, 2026	Press reporting
S40	Huang: Nvidia's investments in OpenAI and Anthropic likely its last	TechCrunch	https://techcrunch.com/2026/03/04/jensen-huang-says-nvidia-is-pulling-back-from-openai-and-anthropic-but-his-explanation-raises-more-questions-than-it-answers/	Mar 4, 2026	Press reporting
S41	Nvidia involved in >\$750B of proposed AI partnerships and financing; circularity warnings	Axios	https://www.axios.com/2026/08/05/nvidias-ai-chips-financing	Aug 5, 2026	Press reporting
S42	AI investment growing ~46% faster than revenue vs 32% in 2001 telecom excess	Allianz Research	https://www.allianz.com/content/dam/onemarketing/azcom/Allianz_com/economic-research/publications/specials/en/2026/march/2026_03_25_AI.pdf	Mar 25, 2026	Analyst estimate
S43	Chip useful life the single most influential variable; ~\$7.6T AI capex 2026–31	Goldman Sachs, "Tracking Trillions"	https://www.goldmansachs.com/insights/articles/tracking-trillions-the-assumptions-shaping-scale-of-the-ai-build-out	May 1, 2026	Analyst estimate
S44	\$2T annual AI revenue needed by 2030; ~\$800B shortfall	Bain & Company	https://www.bain.com/about/media-center/press-releases/20252/\$2-trillion-in-new-revenue-needed-to-fund-ais-scaling-trend--bain--companys-6th-annual-global-technology-report/	Sep 23, 2025	Analyst estimate
S45	Sales restricted for entities >50% owned by companies in unsupported regions such as China	Anthropic	https://anthropic.com/news/updating-restrictions-of-sales-to-unsupported-regions	Sep 4, 2025	Company statement
S46	US export-control directive suspended all foreign-national access to Fable 5 and Mythos 5; concern a jailbreak involving asking the model to read a specific codebase and fix software flaws; Anthropic disabled both models; argued the standard would halt all new frontier deployments	Anthropic	https://www.anthropic.com/news/fable-mythos-access	Jun 12, 2026	Company statement
S47	Mythos 5 returned only for Annex A entities; Fable 5 still restricted	Axios	https://www.axios.com/2026/06/27/commerce-anthropic-mythos-restrictions-lift	Jun 27, 2026	Press reporting

ID	Claim	Source	URL	Published	Type
S48	Controls lifted; access from 2 Jul; commitments to detect risks, collaborate on standards, report malicious activity; approval to serve critical-infrastructure operators	Al Jazeera	https://www.aljazeera.com/economy/2026/7/1/us-lifts-restrictions-on-powerful-ai-models-fable-mythos-anthropic-says	Jul 1, 2026	Press reporting
S49	Directive relied on ECRA "is-informed" authority; letter unpublished	CSIS	https://www.csis.org/analysis/department-commerce-restricted-access-anthropics-latest-models-what-comes-next	Jun 2026	Independent analysis
S50	Amodei: H200 sales to China "a big mistake"	The Register	https://www.theregister.com/2026/01/20/anthropic_nvidia_china/	Jan 20, 2026	Press reporting
S51	Anthropic urged tighter AI Diffusion Rule; Nvidia rebuttal	Maginative	https://www.maginative.com/article/nvidia-and-anthropic-clash-over-ai-chip-export-controls/	May 1, 2025	Press reporting
S52	BIS moved H200-class exports to China to case-by-case review	Federal Register 2026-00789	https://www.federalregister.gov/documents/2026/01/15/2026-00789/revision-to-license-review-policy-for-advanced-computing-commodities	Jan 15, 2026	Regulatory
S53	TSMC top ten customers 78% of 2025 revenue; R.O.C. political risk factor	TSMC Form 20-F (FY2025)	https://investor.tsmc.com/sites/ir/sec-filings/2025_20F%20Report.pdf	Apr 16, 2026	Regulatory filing
S54	Nvidia \$4.5B H20 charge after license requirement	Nvidia Q1 FY2026 8-K Ex. 99.1	https://www.sec.gov/Archives/edgar/data/1045810/000104581025000115/q1fy26pr.htm	May 28, 2025	Regulatory filing
S55	Nvidia FY2026 10-K: revenue concentration (largest direct customers 22% and 14%, per MarketScreener summary); foreclosed from China data-center market; foundry dependence	Nvidia Form 10-K; MarketScreener	https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm	Feb 25, 2026	Regulatory filing
S56	Anthropic will not accede to mass domestic surveillance or fully autonomous weapons uses	Anthropic, statement from Dario Amodei	https://anthropic.com/news/statement-department-of-war	Feb 26, 2026	Company statement
S57	Agencies directed to cease use; OneGov terminated; Hegseth: no military contractor may conduct commercial activity with Anthropic	FedScoop	https://fedscoop.com/anthropic-claude-dod-federal-agency-fallout-trump-hegseth/	27-Feb 28, 2026	Press reporting
S58	Anthropic reads the designation as covering only use as a direct part of Department of War contracts	Anthropic	https://www.anthropic.com/news/where-stand-department-war	Mar 5, 2026	Company statement
S59	N.D. Cal. (Judge Lin) held designation unlawful: First Amendment retaliation, arbitrary and capricious, due process	TechCrunch	https://techcrunch.com/2026/08/28/anthropic-gets-its-first-court-win-over-the-pentagons-supply-chain-risk-label/	Aug 28, 2026	Press reporting

ID	Claim	Source	URL	Published	Type
S60	D.C. Circuit upheld designation 2–1; timeline incl. "all lawful use" demand; Anthropic considering further review	Defense One; ABC News	https://www.defenseone.com/threats/2026/09/anthropic-lawsuit-supply-chain-risk/416252/ ; https://abcnews.com/Business/anthropic-appeals-court-declines-block-pentagon-blacklisting/story?id=136755690	Sep 25, 2026	Press reporting
S61	CDAO contract with ceiling up to \$200M	DefenseScoop	https://www.defensescoop.com/2025/07/14/pentagon-ai-contracts-musk-xai-google-openai-anthropic-cdao/	Jul 14, 2025	Press reporting
S62	GSA OneGov: Claude at \$1 per agency	GSA	https://www.gsa.gov/about-gsa/newsroom/news-releases/gsa-strikes-onegov-deal-with-anthropic-08122025	Aug 12, 2025	Government
S63	Sacks: Anthropic running a regulatory capture strategy	Bloomberg Law	https://news.bloomberglaw.com/artificial-intelligence/anthropic-is-running-a-strategy-based-on-fear-mongering-sacks	Oct 14, 2025	Press reporting
S64	President attacked Amodi and AI guardrails on Truth Social	Quartz	https://qz.com/trump-ai-guardrails-anthropic-dario-amodei-truth-social-091426	Sep 14, 2026	Press reporting
S65	LTBT structure: Class T stock elects phased-in board majority; PBC purpose	Anthropic	https://www.anthropic.com/news/the-long-term-benefit-trust	Sep 19, 2023	Company statement
S66	Narasimhan appointed; Trust-selected directors become board majority (corroborated by R&D World)	Anthropic; R&D World	https://www.anthropic.com/news/narasimhan-board	Apr 14, 2026	Company statement / press reporting
S67	Trustee Cuéllar appointed chief global affairs officer, leaving the Trust	PYMNTS	https://www.pymnts.com/personnel/2026/anthropic-appoints-former-california-supreme-court-justice-as-first-global-affairs-chief/	Aug 4, 2026	Press reporting
S68	Risk management 50%; RSO designated; no central risk function; board risk committee gaps; limited internal audit	SaferAI Frontier Risk Management Tracker	https://tracker.safer-ai.org/company/anthropic/	Jul 1, 2026	Independent rating
S69	RSP version history; v3.4 in force from Jul 8, 2026	Anthropic, Responsible Scaling Policy	https://www.anthropic.com/responsible-scaling-policy	accessed Sep 26, 2026	Company statement
S70	RSP v3 separates unilateral commitments from industry recommendations; Risk Reports every 3–6 months	Anthropic	https://anthropic.com/news/responsible-scaling-policy-v3	Feb 24, 2026	Company statement
S71	RSP v3 removed pause commitment, downgraded SL-4 security to recommendation; external review narrowly triggered	GovAI	https://www.governance.ai/analysis/anthropics-rsp-v3-o-how-it-works-whats-changed-and-some-reflections	Mar 5, 2026	Independent analysis
S72	Aug 2026 Risk Report: misalignment risk raised to "low"; internal "Model 2" not to be released; LTBT external-review power unused	Unite.AI	https://www.unite.ai/anthropic-raises-misalignment-risk-to-low-and-shelves-internal-model-2/	Aug 14, 2026	Press reporting of company disclosure

ID	Claim	Source	URL	Published	Type
S73	RSP noncompliance / anti-retaliation policy is only a partial whistleblowing policy	AI Whistleblower Initiative	https://aiwhistleblowerinitiative.substack.com/p/anthropic-publishes-their-partial	Dec 8, 2025	Independent analysis
S74	ASL-3 protections activated; >100 security controls	Anthropic	https://anthropic.com/news/activating-asl3-protections	May 22, 2025	Company statement
S75	Mythos Preview via Project Glasswing only; \$100M usage credits; not planned for general availability	Simon Willison	https://simonwillison.net/2026/Apr/7/project-glasswing/	Apr 7, 2026	Independent commentary on company disclosure
S76	Anthropic endorses SB 53	Anthropic	https://anthropic.com/news/anthropic-is-endorsing-sb-53	Sep 8, 2025	Company statement
S77	SB 53 signed: safety frameworks, whistleblower protection, critical-incident reporting to Cal OES	TechCrunch	https://techcrunch.com/2025/09/29/california-governor-newsom-signs-landmark-ai-safety-bill-sb-53/	Sep 29, 2025	Press reporting
S78	RAISE Act signed: 72-hour incident reporting; penalties to \$1M/\$3M; Anthropic supported	TechCrunch	https://techcrunch.com/2025/12/20/new-york-governor-kathy-hochul-signs-raise-act-to-regulate-ai-safety	Dec 20, 2025	Press reporting
S79	Amodei: 10-year moratorium on state AI laws "too blunt"	Reuters via Yahoo Finance	https://finance.yahoo.com/news/anthropic-ceo-says-proposed-10-112523320.html	Jun 5, 2025	Press reporting
S80	Anthropic to sign EU GPAI Code of Practice	Anthropic	https://www.anthropic.com/news/eu-code-practice	Jul 21, 2025	Company statement
S81	Federal lobbying \$1.6M in Q1 2026; issues incl. DoD procurement, supply-chain risk, acceptable-use policy, export controls	Axios (LDA filings)	https://www.axios.com/2026/04/21/anthropic-outspends-openai-biggest-lobbying-quarter	Apr 21, 2026	Press reporting of regulatory filing
S82	Federal lobbying \$3.1M in 2025; >\$3.5M in H1 2026	Axios (LDA filings)	https://www.axios.com/2026/07/21/anthropic-ramps-up-lobbying-spending-ai-policy-fights	Jul 21, 2026	Press reporting of regulatory filing
S83	\$20M to Public First Action; platform: frontier transparency, non-preemptive federal framework, chip export controls, bio/cyber rules	Anthropic	https://www.anthropic.com/news/donate-public-first-action	Feb 12, 2026	Company statement
S84	Second \$20M to Public First Action (total \$40M)	Anthropic; Axios	https://www.anthropic.com/news/donation-public-first-action	Jul 21, 2026	Company statement
S85	Bartz: training on lawfully acquired books fair use; pirated library acquisition not; class ~482,460 works; release limited to past acquisition conduct, excludes outputs	Kluwer Copyright Blog	https://legalblogs.wolterskluwer.com/copyright-blog/the-bartz-v-anthropic-settlement-understanding-americas-largest-copyright-settlement/	2026	Legal commentary

ID	Claim	Source	URL	Published	Type
S86	Final approval Jul 20, 2026; ~\$3,000 per work	Authors Guild	https://authorsguild.org/news/court-grants-final-approval-anthropic-copyright-settlement/	Jul 21, 2026	Press reporting
S87	\$300M + \$300M + \$450M paid by Aug 19, 2026 (\$1.05B of \$1.5B); final \$450M plus interest due the earlier of Sep 25, 2027 or 30 days after a "qualifying funding event" (not defined in the filing); first distribution ~\$2,203 per work (Dkt. 680, 688)	Settlement Insight	https://settlementinsight.com/news/anthropic-settlement-first-payments-by-november-15-2203-dollars-per-work-not-3000-and-30-days-to-contest-the-split	Sep 7, 2026	Press reporting citing court documents
S88	~350 rightsholders opted out	OpenClassActions	https://openclassactions.com/news/anthropic-copyright-settlement-final-approval.php	Sep 8, 2026	Press reporting
S89	Concord I: lyric guardrail stipulation (Jan 2025); second amended complaint seeking up to \$150,000 per work	Music Business Worldwide	https://www.musicbusinessworldwide.com/music-publishers-file-amended-lyrics-lawsuit-against-anthropic-just-as-ai-firms-separate-1-5b-piracy-settlement-with-authors-wins-court-approval/	Jul 22, 2026	Press reporting
S90	Concord II: 20,517 + 714 works; names Amodei and Mann; alleges torrenting concealed in discovery	IPWatchdog	https://ipwatchdog.com/2026/01/29/music-publishers-file-new-piracy-suit-against-anthropic-alleging-mass-torrenting-copyrighted-works/	Jan 29, 2026	Press reporting of court filing
S91	Concord II characterized as ~\$3B suit using Bartz evidence	California Courts Newsroom (press digest)	https://newsroom.courts.ca.gov/news/music-publishers-use-evidence-unearthed-bartz-hit-anthropic-3b-suit	Jan 30, 2026	Press reporting
S92	Sony Music Publishing and Warner Chappell sue Anthropic, Amodei, Mann over tens of thousands of works	Music Business Worldwide; TechCrunch	https://techcrunch.com/2026/08/29/sony-music-warner-sue-anthropic-alleging-a-brazen-campaign-of-intellectual-property-theft/	Aug 29, 2026	Press reporting of court filing
S93	Reddit v. Anthropic remanded to state court; claims not preempted	Courthouse News Service	https://www.courthousenews.com/reddit-privacy-case-against-anthropic-kicked-back-to-state-court/	Mar 30, 2026	Press reporting
S94	Reddit contract, interference and unfair-competition claims survive (Sep 19, 2026 ruling; forum to be confirmed)	National Law Review	https://natlawreview.com/article/terms-service-just-got-teeth-reddits-scraping-suit-against-anthropic-survives	Sep 24, 2026	Legal commentary
S95	Class action over Claude Max usage limits	Engadget	https://www.engadget.com/2194626/anthropic-hit-with-lawsuit-over-its-claude-max-usage-limits/	Jun 15, 2026	Press reporting
S96	Pascual v. Anthropic: degraded Pro/Max service class action	Top Class Actions	https://topclassactions.com/lawsuit-settlements/lawsuit-news/anthropic-class-action-alleges-claude-subscribers-paid-for-degraded-ai-service/	Aug 2026	Press reporting
S97	Claude Code source leak (~500,000 lines), March 2026	Michael Best & Friedrich	https://www.michaelbest.com/insights/rethinking-trade-secrets-vs-patents-for-software-and-ai-in-light-of-the-anthropic-102mqd2/	Apr 21, 2026	Legal commentary
S98	No verified chatbot wrongful-death case against Anthropic	AI Lawsuit Tracker; Morgan & Morgan	https://ailawsuittracker.com/ai-chatbot-harm-lawsuits/	May 17, 2026; Aug 17, 2026	Tracker / commentary

ID	Claim	Source	URL	Published	Type
S99	Commercial Terms: IP indemnity (outside cap); liability capped at 12 months' fees; consequential damages excluded	Anthropic Commercial Terms of Service	https://www.anthropic.com/legal/commercial-terms	effective Jun 17, 2025	Company statement
S100	Consumer Terms: cap greater of 6 months' fees or \$100; 18+; no arbitration or class waiver	Anthropic Consumer Terms of Service	https://www.anthropic.com/legal/consumer-terms	effective Oct 8, 2025	Company statement
S101	Insurers balk at AI labs' exposures; Anthropic partly using own funds for settlement (FT)	Silicon UK; Slashdot	https://www.silicon.co.uk/e-innovation/artificial-intelligence/openai-anthropic-insurance-626923	Oct 9, 2025	Press reporting (secondary)
S102	"Vibe hacking" extortion of ≥17 organizations; North Korean IT workers; ransomware	Anthropic threat report	https://www.anthropic.com/news/detecting-counteracting-misuse-aug-2025	Aug 27, 2025	Company statement
S103	State-sponsored espionage campaign; ~30 targets; AI performed 80–90% of work	Anthropic	https://www.anthropic.com/news/disrupting-AI-espionage	Nov 13, 2025	Company statement
S104	House Homeland Security requested Anthropic testimony on the campaign	House Committee on Homeland Security	https://homeland.house.gov/2025/11/26/homeland-republicans-request-anthropic-google-quantum-xchange-testimony-following-report-of-ai-assisted-partially-autonomous-prc-cyber-operation/	Nov 26, 2025	Government
S105	Raine v. OpenAI consolidated in JCCP 5431	Lawsuit Informer	https://lawsuitinformer.com/raine-v-openai-lawsuit	2026	Tracker
S106	Seven suits against OpenAI alleging emotional manipulation and suicide coaching	Social Media Victims Law Center	https://socialmediavictims.org/press-releases/smvlc-tech-justice-law-project-lawsuits-accuse-chatgpt-of-emotional-manipulation-supercharging-ai-delusions-and-acting-as-a-suicide-coach/	Nov 6, 2025	Press release (plaintiffs)
S107	Lyons v. OpenAI: motion to dismiss or stay denied	Courthouse News Service	https://www.courthousenews.com/openai-cant-duck-federal-claims-over-murder-suicide-tied-to-chatgpt/	Apr 13, 2026	Press reporting
S108	Character.AI/Google settled Garcia	CBS News	https://www.cbsnews.com/news/google-settle-lawsuit-florida-teens-suicide-character-ai-chatbot/	Jan 7, 2026	Press reporting
S109	Garcia: court declined at pleading stage to hold chatbot output protected speech	Akin	https://www.akingump.com/en/insights/ai-law-and-regulation-tracker/district-court-denies-first-amendment-free-speech-rights-for-ai-chatbot	May 23, 2025	Legal commentary on court document
S110	FTC 6(b) orders to seven companies on companion chatbots; Anthropic not included	FTC	https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions	Sep 11, 2025	Regulatory
S111	44 attorneys general letter to 13 companies incl. Anthropic	Tennessee Attorney General	https://www.tn.gov/attorneygeneral/news/2025/8/25/pr25-43.html	Aug 25, 2025	Government
S112	42 attorneys general demand safety testing, recall procedures, warnings from 13 companies incl. Anthropic	Pennsylvania Attorney General	https://www.attorneygeneral.gov/taking-action/ag-sunday-leads-coalition-of-42-attorneys-general-in-letter-to-a-i-software-companies-demanding-safeguards-to-protect-vulnerable-residents-from-harmful-interactions-with-bots/	Dec 10, 2025	Government

ID	Claim	Source	URL	Published	Type
S113	California SB 243 companion-chatbot law (disclosure, self-harm protocols, private right of action); 12 states with such laws	Sen. Padilla; MultiState.ai	https://sd18.senate.ca.gov/news/first-nation-ai-chatbot-safeguards-signed-law ; https://www.multistate.ai/updates/vol-105-state-ai-companion-chatbot-laws	Oct 2025; Jun 26, 2026	Government / analyst
S114	Section 230 unlikely to apply where the AI generates the content	Moody's	https://www.moody's.com/web/en/us/insights/insurance/230-immunity-for-AI-chatbot-lawsuits.html	Mar 5, 2026	Analyst commentary
S115	Master Settlement Agreement ~\$206B over first 25 years	GAO-01-851	https://www.gao.gov/products/gao-01-851	2001	Government
S116	US v. Philip Morris: RICO scheme to defraud; affirmed 2009	Public Health Law Center	https://www.publichealthlawcenter.org/sites/default/files/resources/tclc-fs-DOJ-litigation-overview-2015.pdf	2015	Academic
S117	FDA modified-risk orders for IQOS (2020, 2022; renewed 2026)	FDA	https://www.fda.gov/tobacco-products/advertising-and-promotion/philip-morris-products-sa-modified-risk-tobacco-product-mrtp-applications	2026	Government
S118	~60 asbestos trusts, ~\$37B assets	GAO-11-819	https://www.gao.gov/products/gao-11-819	2011	Government
S119	Largest trusts paid ≥\$10.9B through 2008; median ~a quarter of claim value	RAND TR-872	https://www.rand.org/pubs/technical_reports/TR872.html	2010	Academic
S120	Johns-Manville Chapter 11 petition, Aug 26, 1982	In re Johns-Manville Corp., 41 B.R. 926	https://law.justia.com/cases/federal/district-courts/BR/41/926/1552176/	1984	Court document
S121	ASC 450-20-25-2: accrue when probable and reasonably estimable	Deloitte DART Roadmap: Contingencies	https://dart.deloitte.com/USDART/home/codification/liabilities/asc450-10/deloitte-s-roadmap-contingencies-loss-recoveries/chapter-2-loss-contingencies-commitments/2-3-recognition	current	Professional commentary
S122	21 CFR 314.80: 15-day reports for serious unexpected adverse events	eCFR	https://www.ecfr.gov/current/title-21/chapter-I/subchapter-D/part-314/subpart-B/section-314.80	current	Regulatory
S123	GSK \$3B resolution incl. off-label promotion and failure to report safety data	US DOJ	https://www.justice.gov/archives/opa/pr/glaxosmithkline-plead-guilty-and-pay-3-billion-resolve-fraud-allegations-and-failure-report	Jul 2, 2012	Government
S124	Vioxx withdrawn 2004; Merck \$4.85B settlement	NPR	https://www.npr.org/2007/11/10/5470430/timeline-the-rise-and-fall-of-vioxx	Nov 10, 2007	Press reporting
S125	47 U.S.C. §230(c)(1)	Cornell LII	https://www.law.cornell.edu/uscode/text/47/230	current	Statute
S126	Lemmon v. Snap: negligent design claim outside §230	9th Cir. via Justia	https://law.justia.com/cases/federal/appellate-courts/ca9/20-55295/20-55295-2021-05-04.html	May 4, 2021	Court document
S127	Haugen testimony: harms confirmed by internal research	Senate Commerce Committee	https://www.commerce.senate.gov/wp-content/uploads/media/doc/Frances%20Haugen%20Written%20Testimony.pdf	Oct 4, 2021	Government
S128	LA JCCP jury: Meta and YouTube negligent in platform design; \$6M	ABC7 Los Angeles	https://abc7.com/post/los-angeles-social-media-addiction-trial-jury-finds-instagram-youtube-liable-landmark-court-case/18771272/	Mar 2026	Press reporting

ID	Claim	Source	URL	Published	Type
S129	Meta settlement with state AGs: \$12.1B guaranteed, up to \$17.1B; age verification and two-hour daily limit for minors	New York Attorney General	https://ag.ny.gov/press-release/2026/attorney-general-james-secures-171-billion-and-groundbreaking-reforms-meta	Aug 26, 2026	Government
S130	MDL 3047 school-district bellwether settled ~\$27M	Insurance Journal (Reuters)	https://www.insurancejournal.com/news/southeast/2026/06/02/872092.htm	Jun 2, 2026	Press reporting
S131	Issuer-pay conflict; Dodd-Frank §939A and §939G; SEC Office of Credit Ratings	House Financial Services Committee memo	https://docs.house.gov/meetings/BA/BA16/20210721/113951/HHRG-117-BA16-20210721-SD004.pdf	Jul 16, 2021	Government
S132	S&P \$1.375B settlement; model decisions influenced by issuer relationships	US DOJ	https://www.justice.gov/archives/opa/pr/justice-department-and-state-partners-secure-1375-billion-settlement-sp-defrauding-investors	Feb 3, 2015	Government
S133	Moody's ~\$864M settlement; CEO certifications	US DOJ	https://www.justice.gov/archives/opa/pr/justice-department-and-state-partners-secure-nearly-864-million-settlement-moody-s-arising	Jan 13, 2017	Government
S134	Abu Dhabi Commercial Bank v. Morgan Stanley: limited opinion protection for privately distributed ratings	S.D.N.Y. via Justia / CourtListener	https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1:2008cv07508/331416/67/	Sep 2, 2009	Court document
S135	Waste Management v. SAP: fake demonstration environments; >\$100M sought; settled 2010	Computerworld	https://www.computerworld.com/article/1570731/waste-management-sues-sap-over-erp-implementation.html	Mar 20, 2008; May 3, 2010	Press reporting
S136	Revlon SAP go-live disruption; material weakness; class actions	Computer Weekly	https://www.computerweekly.com/news/252464278/SAP-disruption-leads-to-Revlon-class-action-lawsuit	2019	Press reporting
S137	Lidl abandons SAP after ~€500M	Computer Weekly	https://www.computerweekly.com/news/252446965/Lidl-dumps-500m-SAP-project	Aug 2018	Press reporting
S138	SAP cloud terms: liability capped at annual fees; consequential damages excluded; carve-outs	SAP GTC for Cloud Services (GSA version)	https://static.carahsoft.com/concrete/files/5517/0293/2214/SAP_GTC_Cloud_DPA_and_DPA_GSA_Approved_10.3.23.pdf	Oct 3, 2023	Contract
S139	Telecom boom and bust: ~\$700B market value lost; equipment investment ~doubled 1996–2000	Richmond Fed Economic Quarterly	https://www.richmondfed.org/~media/richmondfedorg/publications/research/economic_quarterly/2003/fall/pdf/wolman.pdf	Fall 2003	Academic
S140	97% of fibre lines unused at any moment	NPR	https://www.npr.org/2002/07/22/1147090/telecom-glut	Jul 22, 2002	Press reporting
S141	WorldCom capitalized ~\$3.8B of operating costs	SEC Litigation Release 17588	https://www.sec.gov/enforcement-litigation/litigation-releases/lr-17588	Jun 27, 2002	Regulatory
S142	Ohio Telecom Ass'n v. FCC: broadband an information service; net-neutrality order struck	6th Cir. via Justia	https://law.justia.com/cases/federal/appellate-courts/ca6/24-3449/24-3449-2025-01-02.html	Jan 2, 2025	Court document

ID	Claim	Source	URL	Published	Type
S143	SR 11-7 model risk management (2011): effective challenge, validation, governance; superseded by SR 26-2 on Apr 17, 2026	Federal Reserve	https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm	Apr 4, 2011	Regulatory
S144	Reg YY: independent risk-committee chair; CRO reports to risk committee and CEO	eCFR 12 CFR 252 Subpart D	https://www.ecfr.gov/current/title-12/chapter-II/subchapter-A/part-252/subpart-D	current	Regulatory
S145	2026 stress test; Fed to publish loss models for comment	Federal Reserve	https://www.federalreserve.gov/newsevents/pressreleases/bcreg20260624a.htm	Jun 24, 2026	Government
S146	Kaplan on RSP v3: unilateral commitments did not make sense "if competitors are blazing ahead"	Zvi Mowshowitz	https://thezvi.substack.com/p/anthropic-responsible-scaling-policy	Apr 1, 2026	Commentary (quoting company)
S147	Economic Index "Cadences": Claude Code sessions score higher on autonomy than chat	Anthropic	https://www.anthropic.com/research/economic-index-june-2026-report	Jun 26, 2026	Company statement
S148	MIT NADA: 95% of organizations see no P&L return from GenAI spending	Virtualization Review	https://virtualizationreview.com/articles/2025/08/19/mit-report-finds-most-ai-business-investments-fail-reveals-genai-divide.aspx	Aug 19, 2025	Press reporting of academic study
S149	Anthropic reply: 2023 RSP outdated; RSP changes require board approval	Obsolete (Garrison Lovely)	https://www.obsolete.pub/p/exclusive-anthropic-is-quietly-backpedalling	May 22, 2025	Press reporting (company reply)
S150	SR 26-2 (Apr 17, 2026) supersedes and replaces SR 11-7; risk-based approach tailored to model risk profile	Federal Reserve, SR 26-2	https://www.federalreserve.gov/supervisionreg/srletters/SR2602.htm	Apr 17, 2026	Regulatory
S151	Revised interagency guidance explicitly excludes generative and agentic AI models from scope; effective challenge by individuals with sufficient independence; lighter vendor-model validation	Sullivan & Cromwell memo	https://www.sullcrom.com/insights/memo/2026/April/OCC-Fed-FDIC-Issue-Revised-Guidance-Model-Risk-Management	Apr 2026	Legal commentary on regulatory guidance