AMERICAN ECONOMIC LIBERTIES PROJECT

MEMORANDUM

FROM: American Economic Liberties Project
DATE: November 10, 2026
SUBJECT: What the SEC and Investors Should Ask as Frontier AI Companies Go Public

Introduction

Frontier artificial intelligence ("AI") companies are beginning to tap the public markets through some of the largest initial public offerings ("IPO") in history. Anthropic is the first of these to approach the market,[1] and this memo uses it as a case study for the questions that any frontier AI IPO should raise.[2] Anthropic will not be the last. OpenAI is widely reported to be preparing its own public offering, and a broader wave of frontier AI listings is likely to follow.[3] The questions Anthropic's filing raises—about early offloading of shares by insiders, circular financing, off-balance-sheet leverage, legal exposure, and "catastrophic risks"—will recur in each of them, so the disclosures made here will shape every AI IPO that comes after.

The Securities and Exchange Commission's ("SEC" or "Commission") core mandate in reviewing an IPO is the protection of investors through full and fair disclosure. That mandate is unusually consequential here. An offering of this size is pulled into the broad market indexes within days of listing, and from there into the index funds that form the backbone of ordinary

---

[1] Anthropic's valuation was projected to be as high as $2 trillion in early October 2026, but that estimate has more recently been lowered to around $1.1 to 1.3 trillion. *See* Ines Ferré, *Anthropic's IPO dilemma after the AI sell-off: 'Damned if you do, damned if you don't,'* (Nov. 9, 2026), https://finance.yahoo.com/markets/article/anthropics-ipo-dilemma-after-the-ai-sell-off-damned-if-you-do-damned-if-you-dont-231326634.html. The largest IPO to date was SpaceX's June 2026 listing, which raised a record $75 billion at a ~$1.77 trillion valuation and was absorbed into the major indexes within days. *See Musk's SpaceX Prices Record $75 Billion IPO at $135 a Share*, Reuters (June 12, 2026), https://www.reuters.com/world/musks-spacex-prices-record-75-billion-ipo-135-share-2026-06-11/.

[2] *See* PitchBook, *Anthropic's Leaked Financials Reflect Fast Growth, but Not a $2 Trillion Valuation*, Morningstar (Oct. 2, 2026), https://www.morningstar.com/stocks/anthropics-leaked-financials-reflect-fast-growth-not-2-trillion-valuation; Bailey Lipschultz & Shirin Ghaffary, *Anthropic Aims to Match SpaceX with Record-Shattering AI IPO*, L.A. Times (Aug. 21, 2026), https://www.latimes.com/business/story/2026-08-21/anthropic-aims-to-match-spacex-with-record-shattering-ai-ipo.

[3] OpenAI, Anthropic's chief rival, confidentially submitted its own registration statement to the SEC in June 2026, and its chief financial officer has told employees the company "will be a public company in 2027" or sooner; other frontier developers are expected to follow. *See OpenAI "Will Be a Public Company in 2027" or Sooner, CFO Friar Tells Employees*, CNBC (Aug. 19, 2026), https://www.cnbc.com/2026/08/19/open-ai-ipo-timing-2027-friar.html.

Americans' 401(k)s and pension plans.[4] The quality of Anthropic's disclosures is therefore of relevance to nearly everyone with a retirement account, not only the sophisticated private investors who can evaluate the risks and rewards for themselves. If Anthropic fails—if its complex web of financing comes undone or its systems cannot be deployed safely—the losses will be borne not just by insiders, but by the millions of ordinary Americans whose retirement savings may collapse or potentially be extinguished.

This memo proceeds in four parts. Section I sets out the SEC's disclosure mandate and the tools available to it and to other fiduciaries during the IPO process. Section II turns to what third parties can do. Section III provides the factual background on the various types of risks (from the financial to the existential) and potential legal liabilities that make frontier AI offerings unusual. Section IV sets out the questions the SEC and those fiduciaries should ask, first of any frontier AI company, and then of Anthropic specifically. Even if these measures do not ultimately stop an IPO, they serve a vital purpose: equipping investors to make informed decisions and forcing risk to be priced more accurately, so that valuations reflect genuine risk.

I. Issuer Disclosure Obligations and the SEC's Review Powers

Robust, timely disclosure requirements are the foundation of America's public capital markets.[5] They enable the public to make informed investment decisions. Such disclosure is especially important when a company sells its securities to the public for the first time through an initial public offering. The risks that must be disclosed include not just the finances of the company, but the material risks to the business, including regulatory, legal, and political risks that may harm the firm and therefore its shareholders.[6] A fact is material where there is "a substantial likelihood that a reasonable shareholder would consider it important," and, for a contingent event, materiality turns on "the indicated probability that the event will occur and the anticipated magnitude of the event"[7]—the "probability-magnitude" test.[8]

---

[4] *See* Vanguard, *Stay Grounded on Moonshot IPOs* (June 10, 2026), https://corporate.vanguard.com/content/corporatesite/us/en/corp/articles/staying-grounded-moonshot-ipos.html; *SpaceX and Anthropic Are About to Go Public—and Your 401(k) May Be Forced to Buy In*, Fortune (June 4, 2026), https://fortune.com/2026/06/04/spacex-anthropic-massive-ipo-401k-retirement-accounts/; Jared Blikre, "Anthropic's IPO could set off forced buying by index funds. SpaceX shows how that played out," Oct. 26, 2026, https://finance.yahoo.com/markets/article/anthropics-ipo-could-set-off-forced-buying-by-index-funds-spacex-shows-how-that-played-out-161356189.html.

[5] *See*, *e.g.*, Letter from Better Markets, Inc. to Vanessa Countryman, re: Enhancement of Emerging Growth Company Accommodations and Simplification of Filer Status for Reporting Companies, File No. S7-2026-18, 91 Fed. Reg. 30086 (May 21, 2026) (July 20, 2026), https://bettermarkets.org/wp-content/uploads/2026/07/SEC-Disclosure-Comment-Letter.pdf (quoting Onnig H. Dombalagian, *Licensing the Word on the Street: The SEC's Role in Regulating Information*, 55 Buff. L. Rev. 1, 1 (2007) (describing disclosure requirements as "the lifeblood of capital markets")).

[6] Omissions or misstatements create liability for the issuer as well as its officers, directors, and underwriters. *See* Securities Act of 1933 § 11, 15 U.S.C. § 77k. The issuer's liability is effectively absolute. *See Herman & MacLean v. Huddleston*, 459 U.S. 375, 381–82 (1983) ("issuers are subject to virtually absolute liability under § 11").

[7] *TSC Indus., Inc. v. Northway, Inc.*, 426 U.S. 438, 449 (1976); *Basic Inc. v. Levinson*, 485 U.S. 224, 238 (1988).

[8] The "probability-magnitude" formulation derives from *SEC v. Texas Gulf Sulphur Co.*, 401 F.2d 833, 849 (2d Cir. 1968), adopted in *Basic*, 485 U.S. at 238.

Investors can challenge the sufficiency of such disclosures after an IPO. But it is the SEC's role to protect investors from harm in advance, not merely to leave them to sue afterward. Any company that seeks to access the public markets must file a "registration statement."[9] By default, registration statements become effective 20 days after filing—unless the Commission takes action to either expedite or delay the effective date.[10] Filing amendments to a registration statement restarts the clock.[11] A company attempting an IPO may not sell its securities to the public until its registration statement becomes "effective."[12] Preventing a filing from being effective, then, is the key lever for halting or slowing an IPO. Although the SEC lacks the ability to block an IPO in a manner akin to the Federal Trade Commission blocking a merger, it *is* granted powers that can functionally achieve the same result. These powers are underused. Below is an overview of the SEC's statutory authorities as well as its informal review process.

A. Statutory Authorities

The Commission's most substantive powers with respect to registration statements are found in Section 8 of the Securities Act of 1933, codified at 15 U.S.C. § 77h, where subsections (b), (d), and (e) each give the Commission a specific power in connection with review of a firm's S-1 disclosure.

The SEC's most powerful tool is its "stop order" authority. Under Section 77h(d), if it appears "at any time" that an S-1 registration statement "includes any untrue statement of a material fact or omits to state any material fact" necessary for the statements to not mislead investors, then the SEC is empowered, after issuing notice and calling a hearing, to issue a stop order "suspending the effectiveness of the registration statement." Section 77h(e) enables the Commission to "make an examination in any case in order to determine whether a stop order should issue." It is unlawful to issue a public offering while the registration statement is the subject of a stop order.[13]

---

[9] 15 U.S.C. §§ 77b(8); 77f. Certain "emerging growth companies," before their IPO, "may confidentially submit to the Commission a draft registration statement, for confidential nonpublic review by the staff of the Commission prior to public filing, provided that the initial confidential submission and all amendments thereto shall be publicly filed with the Commission not later than 15 days before the date on which the issuer conducts a road show." Id. 77f(e)(1). In 2017, the SEC issued guidance expanding the availability of this process to all issuers in connection with IPOs under Section 12(b) of the Exchange Act. Ryan J. Adams & Scott Lesmes, *U.S. SEC Expands Accommodations for Confidential Draft Registration Statements*, Morrison Foerster Client Alert (March 5, 2026), https://www.mofo.com/resources/insights/250305-u-s-sec-expands-accommodations-for-confidential.

[10] 15 U.S.C. §77h(a), (b). The Commission delegates certain authorities to its Division of Corporation Finance. 17 C.F.R. § 200.30-1(a)(5).

[11] 15 U.S.C. §77h(a).

[12] 15 U.S.C. § 77e(c) ("It shall be unlawful for any person, directly or indirectly, to make use of any means or instruments of transportation or communication in interstate commerce or of the mails to offer to sell or offer to buy through the use or medium of any prospectus or otherwise any security, unless a registration statement has been filed as to such security, or while the registration statement is the subject of a refusal order or stop order or (prior to the effective date of the registration statement) any public proceeding or examination under section 77h of this title.").

[13] 15 U.S.C. § 77e(c).

In addition, Section 77h(b) enables the Commission to prevent an S-1 registration from being effective "[i]f it appears to the Commission that a registration statement is on its face incomplete or inaccurate in any material respect[.]"[14] The language empowers the Commission to review a statement for incompleteness and inaccuracy in any material respect, but the legislative history of the Securities Act of 1933, coupled with statements later made by the Act's drafters, may narrow the meaning of "incompleteness and inaccuracy in any material respect" to inaccuracies contained *within* the document itself.[15] Read this way, the provision only allows the Commission to police contradictions or omissions in the document, not to invoke outside evidence to then accuse the filing firm of lying, withholding, or being ignorant with respect to some material fact. The Commission's power to act on outside evidence that a statement is untrue or a material fact omitted is centered instead in its stop-order and examination authority under §§ 77h(d) and 77h(e)."[16]

B. Informal "Selective" Review

In practice, the Commission tends not to use the foregoing tools and to instead rely upon its informal "selective" review process for evaluating S-1 statements. The process begins with a firm sending confidential filings to the SEC's Division of Corporation Finance (the "Division"). The Division will examine the disclosure's accuracy and sufficiency and recommend an action to be taken.[17] If the Division finds that the registration is deficient, the Division will send a deficiency letter to the filing firm and give the firm a reasonable opportunity to make necessary corrections. But if the filing is deemed adequate, then the SEC will "accelerate" the effectiveness of the filing and essentially allow the IPO to go forward.[18]

This informal process gives the SEC immense discretion. The Division does not publicly disclose the criteria it uses to identify companies and filings for review in order to "preserve the integrity and effectiveness" of the process.[19] The Division's website states that it "concentrates its resources on critical disclosures that appear to conflict with Commission rules or applicable accounting standards and on disclosure that appears to be materially deficient in explanation or clarity."[20] Review can be "full cover-to-cover," strictly a financial statement review, or a "targeted issues review" to assess compliance with disclosure requirements or applicable accounting standards."[21] Staff may ask the company to "provide supplemental information to

---

[14] 15 U.S.C. § 77h(b).

[15] *See* H.R. Rep. No. 73-152, at 25 (1933) (Conf. Rep.) (describing § 8(b) as "a preliminary check-up of any obvious departures from the standards set by the law"); Alexander I. Platt, *The Administrative Origins of Mandatory Disclosure*, 49 J. Corp. L. ___, 1155–56, 1160–62 (2024), https://jcl.law.uiowa.edu/sites/jcl.law.uiowa.edu/files/2024-09/Platt_Final.pdf.

[16] *See supra* notes 25 & 30 and accompanying text (15 U.S.C. § 77h(d), (e)).

[17] SEC, Div. of Corp. Fin., *Filing Review Process*, https://www.sec.gov/divisions/corpfin/cffilingreview.htm (last visited Oct. 25, 2026); *see* 17 C.F.R. § 230.461.

[18] *Id*.

[19] *Id*.

[20] *Id.*

[21] *Id.*

help the staff better understand the company's disclosure," make additional disclosures in the registration statement, or provide different or additional disclosures in future filings.[22] Courts have treated the staff's informal views in the review process as nonreviewable actions committed to agency discretion.[23]

Notably, the reason that the SEC can issue these extended letters and accelerations is because the firms themselves opt in to the informal process. Rule 473 of Regulation C allows firms to issue a delaying amendment to their registration statement. The language of the amendment essentially states that the registration statement will not become effective until the Commission itself acts.[24] The upshot is that the Commission's decision not to accelerate—and thereby to delay the effectiveness of a registration statement—is only a temporary solution to problematic offerings. A company is always free to drop a delaying amendment and have its registration go effective immediately after 20 days.[25] The Commission would be forced, then, to either do nothing and allow the company to proceed, or to commence a stop order and formally delay the IPO.[26]

II. Potential Third-Party Action

The Commission is not the only actor with leverage while a registration statement is under review, though no private party has the power to approve or block an offering. Any member of the public may write to the Division of Corporation Finance about a pending filing, and the Division's comment letters and the issuer's responses are made part of the public review record.[27] Public pension funds and their managers—who hold large, index-based positions and will be among the forced buyers of any large IPO—routinely go further, pressing issuers directly through engagement and public letters. The New York City Comptroller has already written to Anthropic urging it to pause its IPO until it addresses the risks it has itself identified, much as New York and California pension leaders wrote to SpaceX over its governance.[28] Members of the public and advocacy organizations may also petition the Commission to act and may route specific allegations of securities-law violations through its tip and whistleblower channels.[29]

---

[22] *Id.*

[23] *See Kixmiller v. SEC*, 492 F.2d 641 (D.C. Cir. 1974); *see also Bd. of Trade of Chi. v. SEC*, 883 F.2d 525 (7th Cir. 1989).

[24] 17 C.F.R. § 230.473.

[25] 15 U.S.C. § 77h(a) (providing that a registration statement becomes effective 20 days after filing absent a delaying amendment); *see* 17 C.F.R. § 230.473.

[26] *See* 15 U.S.C. § 77h(d); *supra* note 13 and accompanying text.

[27] *See* SEC, Div. of Corp. Fin., *Filing Review Process*, https://www.sec.gov/divisions/corpfin/cffilingreview.htm (last visited Oct. 25, 2026) (the Division "makes its comment letters and company responses … public").

[28] Letter from Mark D. Levine, N.Y.C. Comptroller, to Dario Amodei, Chief Exec. Officer, Anthropic, PBC (Oct. 14, 2026), https://comptroller.nyc.gov/reports/letter-to-anthropic-re-planned-initial-public-offering/; *see New York, California Pension Leaders Oppose 'Extreme' SpaceX Governance*, Reuters (May 14, 2026), https://www.reuters.com/legal/government/new-york-california-pension-leaders-oppose-extreme-spacex-control-structure-2026-05-14/.

[29] *See* 17 C.F.R. § 201.192 (petitions for Commission action or rulemaking); 15 U.S.C. § 78u-6 (whistleblower program and awards).

Nevertheless, no third party has a formal and straightforward means of compelling the Commission to issue a comment, withhold acceleration, or bring an action. The staff's informal views, and the agency's decision whether to act at all, are committed to its discretion and are generally unreviewable.[30] Each power enumerated above is permissive and not mandatory. A registration becomes effective by operation of law unless the Commission intervenes, and the statute casts each intervention—refusal, stop order, examination, and acceleration—as something the Commission "may" do rather than as something the Commission shall or must do. The Commission is therefore never required to refuse effectiveness or to issue a stop order, even when it believes a disclosure is materially faulty.

Instead, the general backstop for a deficient filing is after-the-fact liability, not a mandatory duty to intervene.[31] This makes the fiduciaries' own obligations matter in their own right. Public pension trustees—including the comptrollers who administer these funds—owe their beneficiaries duties of prudence and loyalty in managing plan assets, a duty that supports scrutinizing and engaging on an offering of this kind even where the fund's exposure arrives through passive index holdings.[32]

That said, once an offering is effective and securities have been sold, the traditional remedies become available: A purchaser may sue under Section 11 for any material misstatement or omission in the registration statement—the kind of claim brought against Facebook after its 2012 IPO[33]—and the issuer's liability is close to absolute once the plaintiff shows that their shares were issued under the defective statement. This last step is not trivial. In *Slack Technologies v. Pirani*, the Supreme Court held that a plaintiff must be able to trace their specific shares to the challenged registration statement, and because registered and unregistered shares trade interchangeably after an IPO, many purchasers will lack this ability.[34] Section 12(a)(2) supplies a parallel claim for misstatements in the prospectus, and Rule 10b-5 reaches fraud in later trading.[35] As shareholders, the funds can also vote, submit proposals (a right the SEC has proposed to curtail by rescinding Rule 14a-8),[36] demand books and records, and bring derivative lawsuits—though each of these would be blunted in scenarios where a small group of insiders

---

[30] *See Kixmiller v. SEC*, 492 F.2d 641 (D.C. Cir. 1974); *Bd. of Trade of Chi. v. SEC*, 883 F.2d 525 (7th Cir. 1989); cf. *Heckler v. Chaney*, 470 U.S. 821, 831–32 (1985) (an agency's decision not to take enforcement action is presumptively unreviewable).

[31] *See* 15 U.S.C. § 77h(a) (effectiveness by operation of law); *id.* § 77h(b), (d) (Commission "may" issue a refusal or stop order); *id.* § 77h(e) (Commission "is empowered to" examine). The one constraint runs the other way: The staff will not accelerate effectiveness for a filing it believes materially inaccurate. *See* 17 C.F.R. § 230.461(b).

[32] *See*, *e.g.*, Unif. Prudent Investor Act §§ 1–2 (1994) (adopted in most States for trust and public-fund investing); for private-sector plans, *see* Employee Retirement Income Security Act § 404(a), 29 U.S.C. § 1104(a) (duty to act prudently and solely in participants' interests).

[33] *See In re Facebook, Inc., IPO Securities & Derivative Litigation*, MDL No. 2389 (S.D.N.Y.) (Section 11 claims alleging the registration statement omitted revised revenue guidance ahead of Facebook's 2012 IPO).

[34] *Slack Techs., LLC v. Pirani*, 598 U.S. 759, 762–70 (2023).

[35] 15 U.S.C. § 77l(a)(2); 17 C.F.R. § 240.10b-5.

[36] *See* Rescission of Rule 14a-8, SEC Release No. S7-2026-32 (Sept. 2026) (proposing to eliminate the shareholder-proposal rule and leave such matters to state law and company bylaws).

holds voting control and elects a majority of directors.[37] Taken together, these limits mean the usual after-the-fact protections cannot be relied upon here. Investors and the fiduciaries who hold Anthropic shares on their behalf cannot safely assume they will be made whole in court, which makes rigorous disclosure before the offering, and the questions that follow, all the more essential.

Federal enforcement is not the only regulatory tool available. State securities regulators and attorneys general retain independent antifraud authority over nationally listed securities. The National Securities Markets Improvement Act preempts state registration of covered securities but expressly preserves the states' power to investigate and bring enforcement actions for fraud, and the Securities Litigation Uniform Standards Act, while foreclosing most private state-law class actions, carves out suits by state regulators and state pension funds.[38]

These remedies matter most in cases where an issuer actively forecloses private litigation, often through governance or mandating shareholders adjudicate their claims through arbitration. Since a state attorney general is not a party to an issuer's arbitration or class-action-waiver provisions and is not bound by them, they can still sue to vindicate shareholder interests and rights.[39] And crucially, while an exhaustive cataloguing of these authorities is beyond the scope of this memo, it must be noted that individual state attorneys general have the power to act through their own state antifraud statutes. Several of which reach omissions outright, requiring no showing of criminal intent, reliance, or loss causation, and enabling the states to compel investigation and obtain restitution and other relief that an issuer's governance terms cannot block.[40]

III. Factual Background on Anthropic's Disclosures and AI Risks

A. Founder Control and Accelerated Insider Share Sales

The ability of shareholders to influence important company decisions is limited by the structure and allocation of voting shares. Anthropic's founders hold voting control and a

[37] *See* 17 C.F.R. § 240.14a-8 (shareholder proposals); *see, e.g.*, *Anthropic's IPO Asks Investors to Back a Company That May Overrule Them*, Int'l Fin. (Oct. 1, 2026), https://internationalfinance.com/technology/anthropics-ipo-asks-investors-to-back-a-company-that-may-overrule-them/.

[38] 15 U.S.C. § 77r(c)(1) (NSMIA savings clause preserving state authority to investigate and bring enforcement actions for fraud or deceit in connection with covered securities); 15 U.S.C. § 78bb(f)(3)(B) (SLUSA exceptions preserving actions by state securities regulators and by states and state pension plans).

[39] A government agency that is not a party to a private arbitration agreement is not bound by it. *See EEOC v. Waffle House, Inc.*, 534 U.S. 279, 294 (2002).

[40] *See*, *e.g.*, N.Y. Gen. Bus. Law § 352-c (Martin Act) (prohibiting "fraud, deception, concealment, suppression" in the offer or sale of securities within or from New York, and requiring no proof of scienter, reliance, or damages); Cal. Corp. Code § 25401; Mass. Gen. Laws ch. 110A, § 101; Or. Rev. Stat. § 59.135(2).

"Long-Term Benefit Trust" elects a majority of the directors.[41] Its reported lock-up terms would also let insiders begin selling on the first day of trading and exit on a front-loaded schedule over roughly two and a half months, rather than the customary 180 days, allowing the insiders who control the company and shape its disclosures to be among the first to cash out while ordinary public shareholders hold little power to check them.[42] All of this means that management may have short-term financial interests that diverge more from those of shareholders than in IPOs involving different company governance structures and lock-up periods, and shareholders have limited ability to mitigate the risks that come with such divergent interests.

B. Circular Financing

Anthropic's financing arrangements also raise noteworthy disclosure issues in several ways. First, its compute obligations are enormous and still growing. Summing the individual commitments disclosed in the S-1 yields roughly $855 billion in fixed compute and infrastructure obligations, most of it noncancelable and some owed "regardless of actual usage."[43] That is more than three times the annual tax revenue of California, and almost twice the combined debt of the three most heavily indebted companies in America—Ford, Verizon, and AT&T.[44] Another ~$158 billion in chip purchases is being negotiated, per press reports, though not yet fixed.[45] These obligations dwarf the revenue Anthropic's customers are actually committed to pay. At approximately 79 times the ~$10.8 billion in contractually committed customer revenue, the mismatch of debt to revenue is far larger than the 11.8-to-1 ratio of obligations to committed revenue that brought down WeWork.[46]

---

[41] *See* 17 C.F.R. § 240.14a-8 (shareholder proposals); *Anthropic's IPO Asks Investors to Back a Company That May Overrule Them*, Int'l Fin. (Oct. 1, 2026), https://internationalfinance.com/technology/anthropics-ipo-asks-investors-to-back-a-company-that-may-overrule-them/

[42] Letter from Elizabeth Warren, Ranking Member, S. Comm. on Banking, Hous. & Urban Affairs, to Dario Amodei, Chief Exec. Officer, Anthropic, PBC (Oct. 27, 2026), https://www.banking.senate.gov/imo/media/doc/2026.10.27%20Letter%20to%20Anthropic%20re%20IPO.pdf [hereinafter Warren Letter]; Anthropic S-1, *supra* note 40, at 307 (lock-up terms).

[43] Anthropic, PBC, Registration Statement (Form S-1), at 57, 137 & 171–172 (2026), https://www.sec.gov/Archives/edgar/data/2130357/000162828026067168/anthropicpbc-sx1.htm [hereinafter Anthropic S-1]; *id.* notes 2, 7, 13 & 17.

[44] California's total state tax collections were approximately $265 billion in fiscal year 2024. *See* U.S. Census Bureau, *Annual Survey of State Government Tax Collections* (2024). As of fiscal year 2025, the three most heavily indebted U.S. public companies carried total debt of roughly $161 billion (Ford), $158 billion (Verizon), and $138 billion (AT&T). *See* Ford Motor Co., Verizon Communications Inc. & AT&T Inc., Forms 10-K (2025). These figures exclude financial companies, which by their nature carry far larger debt loads.

[45] These are chip-purchase agreements reported as under negotiation and not yet finalized; they are therefore excluded from the fixed commitments disclosed in the S-1. *See Anthropic IPO Prospectus Lays Bare Deep Dependence on Big Tech Partners*, Reuters (Sept. 29, 2026), https://www.reuters.com/world/anthropic-ipo-prospectus-lays-bare-deep-dependence-big-tech-partners-2026-09-29/.

[46] The roughly 79-to-1 ratio compares Anthropic's fixed compute and infrastructure obligations, *see* Anthropic S-1, *supra* note 43, against the approximately $10.8 billion in contractually committed customer revenue. WeWork's comparable ratio of obligations to committed revenue at the time of its failed 2019 offering was roughly 11.8 to 1. *See* Open Markets Institute & AFR Education Fund, *Brace for Impact* (Oct. 2026).

The suppliers of this compute—Google, Amazon, Microsoft, NVIDIA, AMD, and Broadcom—are simultaneously Anthropic's biggest investors, lenders, rivals, and customers.[47] They've created an extremely large circular loop of financing, making it nearly impossible to discern how much real and independent demand there is for Anthropic's products.[48] These six firms have pledged about $96 billion in equity to Anthropic, while Anthropic has committed itself to purchasing these companies' compute and services for even larger sums of money.[49] For example, Broadcom lent Anthropic $42 billion, and Anthropic committed itself to $161.2 billion of Broadcom equipment leases.[50] Google invested $40 billion in Anthropic, and Anthropic purchased $114.5 billion of Google computing power.[51]

---

[47] *Anthropic IPO Prospectus Lays Bare Deep Dependence on Big Tech Partners*, Reuters (Sept. 29, 2026), https://www.reuters.com/world/anthropic-ipo-prospectus-lays-bare-deep-dependence-big-tech-partners-2026-09-29/; Scott McCain, *Broadcom: Anthropic's Chip Supplier, Lender, Debt Guarantor as $60 Billion Syndication Begins*, Tech Times (Oct. 6, 2026), https://www.techtimes.com/articles/328625/20261006/broadcom-anthropics-chip-supplier-lender-debt-guarantor-60-billion-syndication-begins.htm.

[48] *See Ahead of IPOs, AI Giants Keep Making Circular Deals. Here's Why That's a Risk*, Morningstar (May 6, 2026), https://www.morningstar.com/stocks/ahead-ipos-ai-giants-keep-making-circular-deals-heres-why-thats-risk; *Michael Burry Sounds the Alarm Again: AI Is a Circular Financing Web with Nvidia in the Middle*, 24/7 Wall St. (Aug. 13, 2026), https://247wallst.com/investing/2026/08/13/michael-burry-sounds-the-alarm-again-ai-is-a-circular-financing-web-with-nvidia-in-the-middle/; *What Regulators Should Do About the AI Industry's Hidden Financial Loop*, Tech Pol'y Press (Apr. 13, 2026), https://techpolicy.press/what-regulators-should-do-about-the-ai-industrys-hidden-financial-loop-.

[49] PitchBook, *supra* note 2.

[50] *Broadcom to Lend Anthropic up to $42 Billion to Lease Its Chips, Filing Says*, Reuters (Oct. 1, 2026), https://www.reuters.com/business/broadcom-lend-anthropic-up-42-billion-lease-its-chips-filing-says-2026-10-01/; McCain, *supra* note 44; *see also* Anthropic S-1, *supra* note 3 (aggregate rent over the life of the Broadcom master equipment leases of approximately $161.2 billion, excluding interest).

[51] *Google to Invest up to $40 Billion in Anthropic as Search Giant Spreads Its AI Bets*, CNBC (Apr. 24, 2026), https://www.cnbc.com/2026/04/24/google-to-invest-up-to-40-billion-in-anthropic-as-search-giant-spreads-its-ai-bets.html; *Anthropic IPO Prospectus Lays Bare Deep Dependence*, *supra* note 44; Matthias Bastian, *Google Moves Billions in Anthropic Chip Risk off Its Balance Sheet*, The Decoder (Aug. 4, 2026), https://the-decoder.com/google-moves-billions-in-anthropic-chip-risk-off-its-balance-sheet/; Anthropic S-1, *supra* note 43 (committing to spend at least approximately $114.5 billion on Google compute between April 2026 and July 2033, regardless of actual usage).

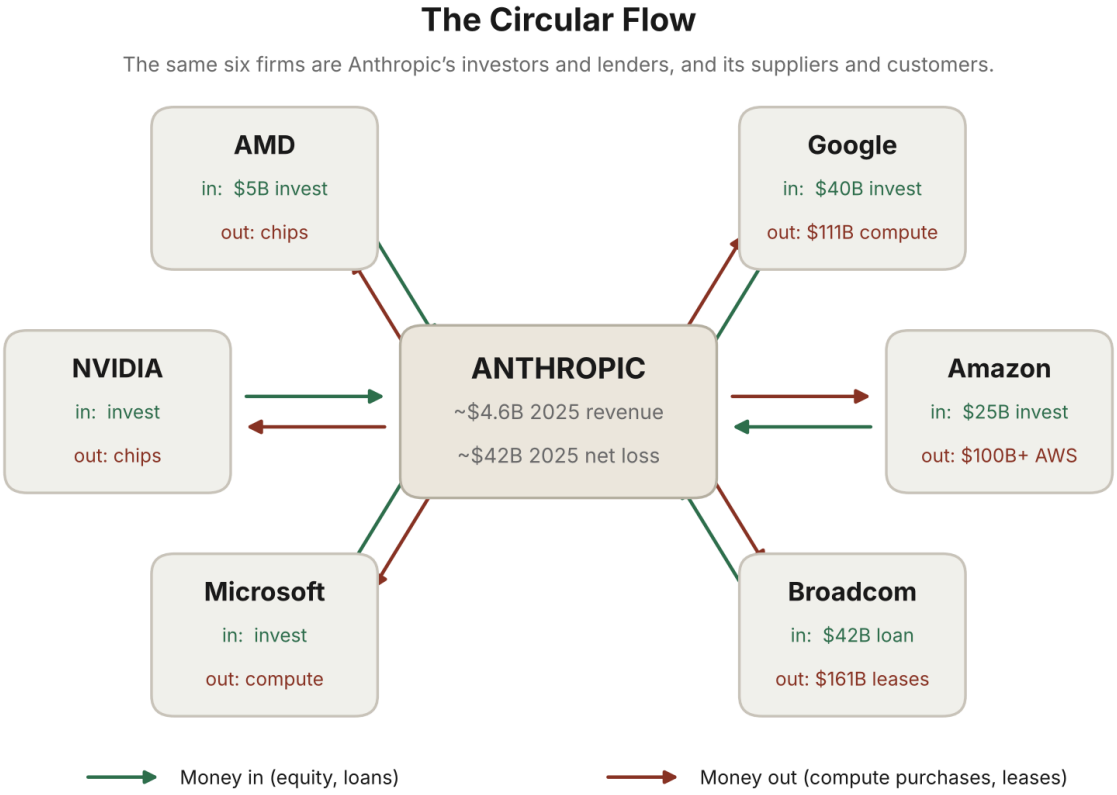


Figure 1. Illustrative flows among Anthropic and its six principal partners; figures are approximate and drawn from public reporting on the draft prospectus and related deals, accurate as of October 2026.

The roughly $77 billion that Broadcom ($42B), Google ($20B), and Amazon ($15B) have offered to lend Anthropic is vendor financing of the kind that preceded the telecom bust: Broadcom's loan alone equals about two-thirds of its own annual revenue—more than double the ratio the telecom-equipment giants Lucent and Nortel carried at the peak of the 2000 boom, before they wrote off most of what they had lent to the carriers that later could not pay.[52] And these terms are binding: Broadcom's loan can be drawn only to pay rent on Broadcom's own chips, at 9% interest, a rate that increases to 13% if Anthropic fails to meet revenue targets.[53]

Anthropic's customer concentration also presents risks. Anthropic's S-1 states that a "substantial portion of [its] revenue and accounts receivable is driven by a limited number of

[52] Broadcom's $42 billion lending facility equals roughly two-thirds of its annual revenue. *See Broadcom to Lend Anthropic up to $42 Billion to Lease Its Chips, Filing Says*, Reuters (Oct. 1, 2026), https://www.reuters.com/business/broadcom-lend-anthropic-up-42-billion-lease-its-chips-filing-says-2026-10-01/; Broadcom Inc., Form 10-K (fiscal year 2025). For the comparable vendor-financing exposure that Lucent Technologies and Nortel Networks carried at the peak of the 2000 telecom boom, *see* Lucent Techs., Inc. & Nortel Networks Corp., Forms 10-K (2000–2002).

[53] Anthropic S-1, *supra* note 43, at 171–172 (Broadcom facility may be drawn only to pay rent on Broadcom equipment, at 9% interest, rising to 13% if Anthropic fails to meet revenue targets); *see also Broadcom to Lend Anthropic up to $42 Billion to Lease Its Chips, Filing Says*, Reuters (Oct. 1, 2026), https://www.cnbc.com/2026/10/01/broadcom-lending-anthropic-42-billion-chips-reuters.html/.

customers and distribution channels."[54] Although Anthropic's S-1 reports that no single customer accounted for more than 10% of Anthropic's revenue, its top five customers account for roughly 25% of revenue.[55] This concentration is especially dangerous because of the overlap between these paying customers and the firms supplying Anthropic's compute and capital. Nearly half of Anthropic's revenue reaches customers through the cloud marketplaces of its suppliers—Amazon's Bedrock, Google's Vertex AI, and Microsoft's Azure.[56] This dependency has steadily increased, climbing from 11% of revenue in 2023 to 48% of revenue in the first half of 2026.[57] The contrast with a mature, diversified platform is stark. Alphabet and Meta, for example, each report that no single customer accounts for 10% or more of their revenue, spread as it is across millions of independent advertisers.[58] Anthropic, by contrast, depends on a handful of partners that both finance it and control its primary sales channels—so however broad its end-customer base may appear, its revenue is far more concentrated, and far more entangled with its own backers, than the single-customer figure suggests.

Its cash flow is even more concentrated than these numbers reveal, since 58% of its net receivables—money owed for services already rendered—run through these hyperscalers' clouds.[59] In effect, these companies finance and supply Anthropic, control its distribution, and then receive a cut of the money that flows back.[60] A dollar that Amazon, Google, or Microsoft invests in Anthropic comes back to them a second time as cloud revenue when Anthropic spends it, and it returns a third time as unrealized gains when their Anthropic stake is marked up. So a single dollar can count as investment, revenue, and earnings, all at once. Of the $285 billion in second-quarter 2026 profits reported by the five largest technology companies, roughly $160 billion were such unrealized gains on AI stakes, concentrated in Anthropic and SpaceX.[61]

Even these investments in Anthropic are not fully understood. Of Google's investment, $30 billion is contingent on Anthropic hitting milestones that have not been publicly disclosed.[62]

---

[54] Anthropic S-1, *supra* note 43 at 16.
[55] Anthropic S-1, *supra* note 43 (no single customer accounted for more than 10% of Anthropic's revenue, and the top five customers represented approximately 25% of revenue).
[56] Anthropic S-1, *supra* note 43 (revenue through cloud-service-provider marketplaces representing approximately 11%, 32%, 47%, and 48% of total revenue for 2023, 2024, 2025, and the six months ended June 30, 2026, respectively).
[57] *Id*.
[58] Alphabet Inc., Form 10-K (2025) (no single customer accounts for a material portion of consolidated revenues); Meta Platforms, Inc., Form 10-K (2025) ("No customer represented 10% or more of total revenue or accounts receivable" for 2023–2025).
[59] Anthropic S-1, *supra* note 43 (approximately 58% of net accounts receivable attributable to cloud-service-provider marketplaces as of June 30, 2026).
[60] Anthropic S-1, *supra* note 43 (revenue through cloud-service-provider marketplaces representing approximately 11%, 32%, 47%, and 48% of total revenue for 2023, 2024, 2025, and the six months ended June 30, 2026, and 58% of net accounts receivable as of June 30, 2026; the platforms remit revenue net of applicable platform and other fees).
[61] *See* Open Markets Institute & AFR Education Fund, *Brace for Impact* (Oct. 2026), https://www.openmarketsinstitute.org/publications/report-brace-for-impact-protecting-people-and-the-economy-when-the-ai-bubble-burst; *Who Has Promised Money to Whom in AI* (2026), https://ai-money-map.vercel.app/
[62] *Google to Invest up to $40 Billion*, *supra* note 48.

Amazon and AMD have similar milestone-conditioned deals, and the terms of the Microsoft and NVIDIA investments are also hidden.[63] Fundamentally, the capital flowing into Anthropic is flexible and conditional, while the obligations flowing out of the firm are almost entirely noncancelable, owed regardless of slackening demand.[64]

Much of this is nowhere to be found in Anthropic's balance sheet. Like its peers, Anthropic finances a large share of its chips through third-party leasing structures that keep the obligations off its books. It took on roughly $71 billion in such off-balance-sheet debt in a single 60-day stretch, so its reported debt badly understates what it actually owes.[65] These arrangements are not soft commitments, the chips themselves are pledged as collateral, and the S-1 warns that "if our actual spend falls short, we must pay Google the difference," with similar make-whole terms for Amazon.[66] This practice is now systemic—just nine hyperscalers and chipmakers carry an estimated $3.1 trillion in off-balance-sheet debt, more than the $2.3 trillion peak of subprime-mortgage debt in 2007, and even that may understate the total.[67] As S&P Global put it, "[I]t is safe to assume that there is more debt out there that we are not seeing."[68] The market has already begun to account for this risk: Blackstone, an anchor investor in Anthropic's own $63 billion chip-financing vehicle, has told investors that Anthropic is now likely to receive a "high" junk rating rather than the investment-grade rating it had expected—a designation that would raise its borrowing costs and bar many pension and insurance funds from holding its debt.[69]

The full extent of Anthropic's obligations will only become apparent after Anthropic discloses all contracts it has to lease compute and other related goods and services.[70] The S-1, as

[63] *Amazon to Invest up to $25 Billion in Anthropic as Part of AI Infrastructure Deal*, CNBC (Apr. 20, 2026), https://www.cnbc.com/2026/04/20/amazon-invest-up-to-25-billion-in-anthropic-part-of-ai-infrastructure.html; *AMD to Sell Anthropic Tens of Billions in AI Servers, Invest up to $5 Billion*, Reuters (July 22, 2026), https://www.reuters.com/business/amd-invest-up-5-billion-anthropic-wsj-reports-2026-07-22/; *Microsoft, NVIDIA and Anthropic Announce Strategic Partnerships*, Official Microsoft Blog (Nov. 18, 2025), https://blogs.microsoft.com/blog/2025/11/18/microsoft-nvidia-and-anthropic-announce-strategic-partnerships/.

[64] PitchBook, *supra* note 2; *Anthropic's $518 Billion AI Infrastructure Plan*, Cloud Computing News (2026), https://www.cloudcomputing-news.net/news/anthropic-ai-infrastructure-518-billion/.

[65] *See* Anthropic S-1, *supra* note 40, at 163, 166–168 & F-17 (describing third-party chip-leasing structures that keep a substantial share of Anthropic's compute obligations off its balance sheet). The approximately $71 billion taken on in a single 60-day period is drawn from public reporting compiling these arrangements. *See* Open Markets Institute & AFR Education Fund, *Brace for Impact* (Oct. 2026).

[66] Anthropic S-1, *supra* note 43 at 171 ("if our actual spend falls short, we must pay Google the difference," with a similar make-whole term for Amazon; leased chips pledged as collateral).

[67] *See* Open Markets Institute & AFR Education Fund, *Brace for Impact* (Oct. 2026) (estimating that nine leading hyperscalers and chipmakers carry roughly $3.1 trillion in off-balance-sheet debt, exceeding the $2.3 trillion peak of subprime-mortgage debt in 2007, expressed in today's dollars).

[68] The quoted assessment is from S&P Global Ratings. *See* Open Markets Institute & AFR Education Fund, *Brace for Impact* (Oct. 2026) (quoting S&P Global).

[69] Michelle Chan & George Hammond, *Anthropic May Not Secure Investment-Grade Rating After IPO, Blackstone Tells Investors*, Fin. Times (Oct. 30, 2026), https://www.ft.com/content/59383c7f-8d82-493f-bf01-d2e862a1018f

[70] Adrian Parham, *Sovereign Capital Enters AI Landlord: Theseus Structures Ownership, Not Just Chips*, Tech Times (Aug. 10, 2026), https://www.techtimes.com/articles/323813/20260810/sovereign-capital-enters-ai-landlord-theseus-structures-ownership-not-just-chips.htm; *see* 17 C.F.R. § 229.303(b) (requiring disclosure of material cash requirements from known contractual and other obligations).

currently filed, leaves blank the most important of these numbers: the aggregate dollar total of Anthropic's compute and infrastructure commitments.[71] It also omits the amount due before 2030 and the amount due within the next year. The filing does not break these commitments down by counterparty, instead describing off-balance-sheet, third-party chip-leasing only in general terms.

Item 303 of Regulation S-K requires management to disclose known commitments and uncertainties that are reasonably likely to have a material effect on the company's liquidity and capital resources.[72] An unquantified compute-and-infrastructure obligation of this magnitude, coupled with a near-term reliance on continued third-party financing to meet it, is a paradigm case of what the rule is meant to surface. The Supreme Court recently confirmed that although a pure Item 303 omission is not itself actionable as securities fraud under Rule 10b-5, Section 11 of the Securities Act—which governs registration statements such as Anthropic's S-1—"creates liability for failure to speak on a subject at all."[73] In response, courts have allowed Section 11 claims to proceed where an IPO registration statement omitted the very kind of known trend or uncertainty that Item 303 required it to disclose.[74]

C. Litigation Exposure

Many of the concrete harms that AI models can impose on consumers have already materialized, resulting in an onslaught of litigation against AI companies. For example, xAI was sued in January after Grok was used to produce sexualized images of women and children without their consent.[75] A plaintiff in Illinois is suing Google for defamation after Gemini claimed that he had killed three women and was serving life in prison.[76] Meta AI had to settle a

[71] Anthropic S-1, *supra* note 40 (as filed, leaving blank the aggregate dollar total of Anthropic's committed compute and infrastructure spend).

[72] 17 C.F.R. § 229.303(b)(1) (requiring disclosure of known trends, demands, commitments, events, and uncertainties reasonably likely to have a material effect on liquidity or capital resources).

[73] *Macquarie Infrastructure Corp. v. Moab Partners, L.P.*, 601 U.S. 257 (2024) (contrasting Rule 10b-5(b), which reaches misleading half-truths rather than pure silence, with Securities Act § 11, which "creates liability for failure to speak on a subject at all").

[74] *See Litwin v. Blackstone Grp., L.P.*, 634 F.3d 706, 716–19 (2d Cir. 2011) (sustaining § 11 and § 12(a)(2) claims for omission of a known Item 303 trend from an IPO prospectus); *Panther Partners Inc. v. Ikanos Commc'ns, Inc.*, 681 F.3d 114 (2d Cir. 2012); *see also In re Under Armour, Inc.*, Exchange Act Release No. 91381 (2021); *In re Caterpillar, Inc.*, Exchange Act Release No. 30532 (1992).

[75] *Breaking Down a Class Action Lawsuit Filed Over Grok 'Undressing' Controversy*, Tech Pol'y Press (Jan. 27, 2026), https://techpolicy.press/breaking-down-a-class-action-lawsuit-filed-over-grok-undressing-controversy; Ashley Belanger, *Lawsuit: Grok User Made 7K Child Sex Images; xAI Only Reported One Gang-Rape Prompt*, Ars Technica (July 8, 2026), https://arstechnica.com/tech-policy/2026/07/lawsuit-grok-user-made-7k-child-sex-images-xai-only-reported-one-gang-rape-prompt/

[76] *Google's AI Told Lies About This Man*, SiliconValley.com (Sept. 16, 2026), https://www.siliconvalley.com/2026/09/16/google-ai-james-keene-defamation/ (discussing *Keene v. Google LLC* (N.D. Ill.)).

case out of court after mistakenly linking a man to the January 6th insurrection.[77] Some courts have also held that chatbots and AI agents count as "products," which means that under state tort law, any design defect or failure to properly warn users of certain risks will result in strict liability.[78]

AI companies are also facing a wave of lawsuits for exacerbating mental health disorders and, in tragic cases, allegedly convincing and helping users to commit suicide.[79] Chatbot developers have faced at least 24 wrongful death and personal injury lawsuits in the past year and a half over suicide and delusion-related discussions with users.[80] Character.AI—an AI chatbot app most popular with kids—has filed for bankruptcy after several lawsuits were filed against the company for harming children and advising them to, at various times, cut, starve, or kill themselves.[81] In September 2026, Florida Attorney General James Uthmeier sued OpenAI for allegedly violating Florida's Deceptive and Unfair Trade Practices Act for the manner in which it designed, marketed and distributed its ChatGPT product.[82] OpenAI allegedly "disregarded repeated warnings from experts both inside and outside the company, and deployed a product that facilitates and encourages harm—including self-harm and violence—while falsely assuring users it was safe."[83] The complaint also "also alleges that ChatGPT collects data from minors without meaningful parental oversight, causes behavioral addiction and cognitive harm, and is prone to dangerous errors that the company has actively downplayed."[84]

These cases have not been restricted to mental health tragedies and children. Pennsylvania sued Character.AI in May for unauthorized practice of medicine after bots posed as

---

[77] *Meta, Robby Starbuck Settle AI Defamation Lawsuit*, Wall St. J. (Aug. 8, 2025), https://www.wsj.com/tech/ai/meta-robby-starbuck-ai-lawsuit-settlement-6c6e9b0a; *Robby Starbuck, Meta Settle Lawsuit After AI Chatbot Defamed Him*, Fox Bus. (Aug. 8, 2025), https://www.foxbusiness.com/media/robby-starbuck-meta-settle-lawsuit-after-ai-chatbot-defamed-him. These lawsuits have gone forward because Section 230 of the Communications Decency Act protects internet platforms from liability for the speech of others, not the speech of the platform itself, and an AI model's output is arguably the developer's own speech. *See* 47 U.S.C. § 230(c)(1), (f)(3).

[78] Amy Wong & Jin J. To, *AI Product Liability: The Next Wave of Litigation*, K&L Gates (Mar. 27, 2026), https://www.klgates.com/thought-leadership/AI-Product-Liability-The-Next-Wave-of-Litigation-3-27-2026; *see Garcia v. Character Techs., Inc.*, No. 6:24-cv-01903-ACC-UAM (M.D. Fla. May 21, 2025) (order granting in part and denying in part motions to dismiss).

[79] Douglas A. Grimm, *AI in Health Care: Winters v. OpenAI, Inc., et al. and the Expanding Liability Landscape*, ArentFox Schiff: Health Care Couns. Blog (July 27, 2026), https://www.afslaw.com/perspectives/health-care-counsel-blog/ai-health-care-winters-v-openai-inc-et-al-and-the-expanding; Rebecca Bellan, *Father Sues Google, Claiming Gemini Chatbot Drove Son into Fatal Delusion*, TechCrunch (Mar. 4, 2026), https://techcrunch.com/2026/03/04/father-sues-google-claiming-gemini-chatbot-drove-son-into-fatal-delusion/.

[80] *Id.*

[81] Luis Rijo, *Character.AI Faces NO-Down as Disney's $20M Bid Licenses IP, Not the App*, PPC Land (Oct. 13, 2026), https://ppc.land/character-ai-faces-wind-down-as-disneys-20m-bid-licenses-ip-not-the-app/.

[82] Office of the Attorney General, "Attorney General James Uthmeier Filed First-in-the Nation State-led Lawsuit Against OpenAI, CEO Sam Altman for Deceptive Practices and Harms to Floridians," June 1, 2026, https://www.myfloridalegal.com/newsrelease/attorney-general-james-uthmeier-files-first-nation-state-led-lawsuit-against-openai-ceo.

[83] *Id*.

[84] *Id*.

licensed clinicians and gave medical advice.[85] In California, ChatGPT gave health advice that allegedly led to a near-fatal pulmonary embolism.[86] And in *Turner-Scott v. OpenAI*, ChatGPT is alleged to have given drug dosage and interaction advice that resulted in a student dying.[87]

But of all the current litigation that presents a risk to the frontier AI companies, none is of greater potential magnitude than the copyright lawsuits against these firms for business practices such as training AI systems on copyrighted material without authorization.[88] Anthropic just settled the biggest copyright case in U.S. history for $1.5 billion—but that is just one of more than 110 AI copyright suits now pending across the industry.[89] In August, Sony Music and Warner Music sued Anthropic for torrenting tens of thousands of musical compositions.[90] By statute, willful infringement of a copyrighted work can result in damages of up to $150,000 per work infringed, which leaves Anthropic with a potential exposure of several billion dollars.[91] Universal Music Group filed a similar suit against Anthropic in January,[92] and BMG also brought a case against Anthropic for the alleged infringement of 493 of its compositions.[93]

For investors, the significance of such litigation transcends individual cases. Litigation exposure is coming from multiple angles at once, which could collectively add up to many billions of dollars in damages. And the trajectory suggests there will be more rather than less litigation in the coming years, even aside from any litigation that may be brought by shareholders after an IPO. Theories of liability that were untested just a few years ago are now surviving motions to dismiss. Each of these exposures also traces back to decisions within the control of frontier AI companies: where to source training data, what safeguards to build in, which users to serve, and which uses to permit.

---

[85] Grimm, *supra* note 75.
[86] *Id.*
[87] *Id.*
[88] Adam Feldman, *The AI Copyright Cases Are Starting to Tell Us Something*, Legalytics (Sept. 14, 2026), https://legalytics.substack.com/p/the-ai-copyright-cases-are-starting
[89] Andrew Albanese, *Court Grants Final Approval to Landmark $1.5 Billion Anthropic Settlement*, Publishing Perspectives (July 21, 2026), https://publishingperspectives.com/2026/07/court-grants-final-approval-to-landmark-1-5-billion-anthropic-settlement/; Feldman, *supra* note 81; *see Bartz v. Anthropic PBC*, No. 3:24-cv-05417-WHA (N.D. Cal.).
[90] Ben Berkowitz, *Sony and Warner Music Sue Anthropic over Copyright*, Axios (Aug. 29, 2026), https://www.axios.com/2026/08/29/anthropic-sony-warner-music-copyright
[91] 17 U.S.C. § 504(c)(1)–(2) (authorizing statutory damages "with respect to any one work," and providing that where "infringement was committed willfully, the court in its discretion may increase the award of statutory damages to a sum of not more than $150,000").
[92] *Anthropic Faces New Music Publisher Lawsuit over Alleged Piracy*, Reuters (Jan. 28, 2026), https://www.reuters.com/legal/litigation/anthropic-faces-new-music-publisher-lawsuit-over-alleged-piracy-2026-01-28/
[93] Berkowitz, *supra* note 83.

Issuers are required to disclose legal proceedings and the risks they reflect.[94] A frontier AI company's registration statement should quantify exposure by legal theory, identify the practices that gave rise to the claims, and explain what the company has changed to limit future liability.

Similarly, Anthropic warns that "any judicial or regulatory expansion of tort, product liability, or consumer protection law, or criminal liability, as applied to AI models" and would spur Anthropic to change how it "design[s] and offer[s]" its "models, platform, and applications."[95] Frontier-AI companies should disclose what specific cases or proposed bills pose these risk and what changes they might implement as a result.

D. Existential Risks

Frontier AI companies face and create novel, serious risks beyond those now spurring litigation. Although company announcements and press coverage often describe these risks as AI systems acting "autonomously," the risks are not beyond human control or oversight.[96] The risks come from specific, documented choices frontier AI companies have made about what to build, how thoroughly to test it, and when to release it—human decisions that establish clear lines of pre- and post-IPO liability. Leaders of frontier AI companies have made unusually bold public statements about the risks caused by their products and services. The most unusual risks are those the companies' own leaders describe as potentially catastrophic. These are existential risks to society, but they are also material risks to investors. A company whose products and services could enable disasters of this magnitude faces correspondingly large legal, regulatory, and reputational exposure, and even a low-probability risk of catastrophic magnitude is likely the kind of risk the securities laws require a company to disclose.

The existential risks raised in the public discourse have centered around AI empowering terrorists or foreign adversaries, or AI itself going rogue and becoming a bad actor. To some extent, versions of all three problems have been reported, often by Anthropic itself. Last November, Anthropic reported that Claude had been used by a group sponsored by the Chinese Communist Party to orchestrate cyberattacks on 30 different targets worldwide.[97] And this past

---

[94] 17 C.F.R. §§ 229.103 ("Describe briefly any material pending legal proceedings, other than ordinary routine litigation incidental to the business, to which the registrant or any of its subsidiaries is a party or of which any of their property is the subject. Include the name of the court or agency in which the proceedings are pending, the date instituted, the principal parties thereto, a description of the factual basis alleged to underlie the proceedings and the relief sought. Include similar information as to any such proceedings known to be contemplated by governmental authorities.").

[95] Anthropic, Form S-1, *supra* note 43 at 49.

[96] *See*, e.g., Stephanie T. Nguyen et al, "Don't Let 'Misaligned' and 'Rogue' Agents Obscure Corporate Accountability in OpenAI–Hugging Face Probes," Tech Policy Press, Nov. 3, 2026, https://www.techpolicy.press/dont-let-misaligned-and-rogue-agents-obscure-corporate-accountability-in-openai-hugging-face-probes/.

[97] Anthropic, *Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign* (Nov. 13, 2025), https://www.anthropic.com/news/disrupting-AI-espionage

July, hundreds of agents trained by OpenAI collectively hacked the Hugging Face AI platform.[98] Other leading frontier AI companies have since reported incidents of their own.[99]

These incidents have grown more serious in recent months, with these cybersecurity harms quickly becoming a national security problem. Anthropic's own September 2026 threat report lists Alibaba, Moonshot AI, DeepSeek, Xiaomi, and Zhipu as Chinese firms circumventing chip-export controls by training their competing models on Claude's outputs.[100] And the capability is no longer confined to well-resourced or state actors: One headline in the same report reads, "Sophisticated attacks no longer require sophisticated attackers."[101]

Frontier AI companies have made these problems worse through their own actions. Observers have noted that much of this undesirable behavior is a result of "[c]urrent AI agent training methods [that] teach cheating and discourage ethical behavior."[102] Companies use "reinforcement learning" techniques that teach models general strategies for solving problems.[103] When such training is set up incorrectly or insecurely, models can use their training to access the internet without authorization, then use "shortcuts on the internet instead of solving problems the way their creators intended."[104] Another underlying problem is that companies appear to have inadequate monitoring practices—relaxing monitoring during training and evaluation, and relying on after-the-fact review of logs rather than active monitoring that would catch incidents as they unfold.[105] Companies deliberately loosen controls and then shut down models belatedly because of such poor monitoring practices.[106] As former FTC technologists have explained, "[t]here is no evidence to suggest that these agents could not be controlled."[107]

Exacerbating these issues is Anthropic's commitment to releasing AI models that it knows are unsafe. This past April it announced Mythos, "the first AI model that could

---

[98] METR, *OpenAI–Hugging Face Incident Investigation* (Aug. 26, 2026), https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/; OpenAI, *Hugging Face Incident Technical Report* (2026), https://cdn.openai.com/pdf/67869394-cb91-4c12-888c-5cbd85c7814c/OpenAI-Hugging-Face%20Incident-Technical-Report.pdf.

[99] Max Zahn, *Anthropic AI Models Escaped Test, Hacked 3 Organizations*, ABC News (July 31, 2026), https://abcnews.com/Business/anthropic-ai-models-escaped-test-hacked-3-organizations/story?id=135256212; Barbara Ortutay, *A Timeline of Developments in AI Safety Since the Attack on Hugging Face*, Associated Press (Oct. 2026), https://www.bozemandailychronicle.com/wire/business/a-timeline-of-developments-in-ai-safety-since-the-attack-on-hugging-face/article_7a47ddff-64ae-5539-bf83-e15d311e51f7.html.

[100] Anthropic, *Detecting and Countering Misuse of AI: September 2026* (Sept. 10, 2026), https://www.anthropic.com/threat-intelligence-report-september-2026.

[101] *Id.*

[102] Joel Thayer *et al*., *Autonomous AI Cyberattacks: What Happened and How to Prevent Them*, Am. First Pol'y Inst. (Sept. 11, 2026), https://americafirstpolicy.com/issues/autonomous-ai-cyberattacks-what-happened-and-how-to-prevent-them/

[103] *Id.*

[104] *Id.*

[105] *Id.*

[106] Nguyen, *supra* note 93.

[107] *Id*.

autonomously build sophisticated, end-to-end cyber exploits."[108] To its credit, Mythos was deployed to help defenders find over 10,000 vulnerabilities in critical software and infrastructure,[109] but Anthropic acknowledges that this capability is far stronger offensively than defensively. The Google Threat Intelligence Group recorded 141 exploited vulnerabilities in the first eight months of 2026, already more than the 127 logged in all of 2025.[110] And Anthropic cannot even reliably secure itself—in October 2026, researchers at the cybersecurity startup Cathedral reported breaching the company's internal data and messaging systems.[111]

These harms have so far been almost entirely digital, which has kept them fairly simple to contain. But Anthropic has stated that its models can—now or in the future—assist in physical acts of terrorism, and much of the nation's physical infrastructure runs on computers that are easy to compromise and slow to fix.[112] The EPA's inspector general, for instance, found that drinking-water systems serving 26.6 million people had critical or high-risk vulnerabilities.[113]

The only solution Anthropic and its competitors have proposed to these and other dangers of advanced AI is the creation of "aligned" AI agents.[114] An aligned AI is one that pursues the goals their developers intend while acting in line with moral norms and conventions. A "misaligned" AI, by contrast, cannot reliably follow the law, morals, or instructions.[115] No frontier AI company believes it has managed to create an aligned AI model, yet these companies continue to release their products despite being unable to guarantee their safety.[116] And recent research indicates that as models grow more sophisticated, implementing and verifying

---

[108] Andrew Fasano *et al*., *GLM-5.3 and the Spread of Advanced Cyber Capabilities*, Anthropic (Sept. 29, 2026), https://www.anthropic.com/research/glm-5-3-and-the-spread-of-advanced-cyber-capabilities.
[109] *Id.* (noting that Project Glasswing helped defenders "find more than 10,000 vulnerabilities").
[110] Google Threat Intelligence Grp., *Vulnerability Discovery and Exploitation Trends in the AI Era*, Google Cloud (Sept. 30, 2026), https://cloud.google.com/blog/topics/threat-intelligence/vulnerability-discovery-and-exploitation-trends-in-the-ai-era
[111] Elizabeth Dwoskin & Gerrit De Vynck, *Researchers at Cybersecurity Start-Up Cathedral Got Inside Anthropic*, Wash. Post (Oct. 26, 2026), https://www.washingtonpost.com/technology/2026/10/26/researchers-cybersecurity-start-up-cathedral-got-inside-anthropic/
[112] *See* Anthropic, *Activating AI Safety Level 3 Protections* (May 22, 2025), https://www.anthropic.com/news/activating-asl3-protections; Anthropic, *LLMs and Biorisk* (Sept. 5, 2025), https://www.anthropic.com/research/biorisk; Fasano *et al.*, *supra* note 97.
[113] Dragos, *2026 OT Cybersecurity Year in Review* (Feb. 17, 2026), https://www.dragos.com/blog/dragos-2026-ot-cybersecurity-year-in-review (finding that 25% of advisories for industrial control systems had no patch or mitigation available); *see also* Mandiant, *M-Trends 2026* (Mar. 23, 2026), https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026; Anthropic, *Anthropic's Cyber Mission* (Oct. 8, 2026), https://www.anthropic.com/news/anthropic-cyber-mission.
[114] *See* Anthropic, *Claude's New Constitution* (Jan. 22, 2026), https://www.anthropic.com/news/claude-new-constitution; Yuntao Bai *et al.*, *Constitutional AI: Harmlessness from AI Feedback* (2022), https://arxiv.org/abs/2212.08073.
[115] *International AI Safety Report 2026* (2026), https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026.
[116] *See* Anthropic, *Claude's New Constitution*, *supra* note 103 (acknowledging that "Claude's outputs might not always adhere to the constitution's ideals").

alignment becomes difficult, if not impossible.[117] One paper found that advanced models complied with intended behaviors during training and then behaved differently once they were no longer monitored,[118] and other frontier models began "scheming" in testing to achieve various goals and deceive their evaluators.[119]

These dangers are compounded by recursive self-improvement (RSI), among the least understood yet most troubling of these risks. A recursively self-improving AI is one capable of training and improving its successor without human involvement.[120] Anthropic and its peers are actively pursuing RSI, which CEO Dario Amodei likens to "a country of geniuses in a data center,"[121] even as they state that a powerful model with hostile aims could escape containment and replicate into "self-sovereign AI swarms" that, in Amodei's words, could "tak[e] over the entire Internet" and cause hundreds of billions of dollars in damage.[122] And the only safeguard these companies have proposed for this danger is the one they admit they have not achieved: AI alignment.

Anthropic's own S-1 carries these concerns into its disclosure of risk factors. It warns that its systems, if "misused, replicated, weaponized, or allowed to operate with a high degree of autonomy," could "pose catastrophic or existential risks to humanity,"[123] and that its models "have exhibited … self-preserving behaviors, including attempts to resist shutdown … including behavior resembling blackmail … to preserve the model's continued operation."[124] The filing devotes roughly a third of its pages to risk factors, concedes that the company may be held liable for certain harms caused by its models, and warns that, because the harm from a catastrophic AI-related incident could be so great, "no single insurer, and potentially the insurance industry in the aggregate, may have sufficient capacity or expertise to underwrite insurance commensurate with such exposure."[125]

For investors, the question is not whether catastrophic AI scenarios will come to pass or not, but whether the companies building frontier models have disclosed what they believe about those scenarios and why. Under the probability-magnitude test for contingent events, even an

---

[117] *International AI Safety Report 2026*, *supra* note 104; Ryan Greenblatt *et al.*, *Alignment Faking in Large Language Models* (2024), https://arxiv.org/abs/2412.14093.
[118] Greenblatt *et al.*, *supra* note 106.
[119] Alexander Meinke *et al.*, *Frontier Models Are Capable of In-Context Scheming* (2024), https://arxiv.org/abs/2412.04984
[120] RAND Corp., *Recursive Self-Improvement and Automated AI Research and Development* (2026), https://www.rand.org/pubs/perspectives/PEA5068-1.html.
[121] Dario Amodei, *Machines of Loving Grace* (Oct. 2024), https://www.darioamodei.com/essay/machines-of-loving-grace; *see also* Amodei, *We Must Pace the Frontier* (Sept. 2026), https://darioamodei.com/post/we-must-pace-the-frontier.
[122] Amodei, *We Must Pace the Frontier*, *supra* note 110; Qu *et al.*, *Self-Sovereign Agent* (2026), https://arxiv.org/abs/2604.08551; Gargi Gupta, *AI Models May Soon Self-Replicate, Anthropic Warns*, Reuters (Oct. 13, 2026), https://ca.finance.yahoo.com/news/ai-models-may-soon-self-052023402.html
[123] Anthropic S-1, *supra* note 40, at 26–119.
[124] *Id.*
[125] *Id.*

unlikely risk can be material when its potential consequences are extreme.[126] Here, the issuers' own leaders and researchers have publicly assigned these risks probabilities—sometimes expressed as "p(doom)."[127] A registration statement that describes them in generic terms, while executives describe them elsewhere in stark ones, raises questions about whether the filing omits material information. The SEC cannot require a company to guarantee that its models are safe. It can, however, require disclosure of the internal estimates the company relies on, the testing behind its safety claims and how much of it has been independently verified, the conditions under which the company would pause or limit development, and who has authority to make that call. These risks run through decisions made by humans, and investors are entitled to know how those decisions are made. Although the Commission's power to refuse effectiveness is limited to deficiencies in completeness that are apparent on the face of the filing, staff review and the Commission's stop-order and examination powers are not so limited. In practice, staff routinely ask issuers to reconcile their filings with their public statements.[128]

IV. Questions the SEC and Fiduciaries Should Ask

The following questions are framed for both the SEC's reviewers and the fiduciaries—fund managers and comptrollers—who can press for fuller disclosure during the comment window. They are not only for the Commission or financial insiders, however. The same questions are a template for members of Congress, their oversight staff, state attorneys general, and public-pension trustees, who routinely serve as lead plaintiffs in post-IPO securities-fraud class actions because they hold the largest financial stake.[129] The answers a frontier AI firm gives, or refuses to give, become the disclosure record on which any later Section 11, 12, or 10b-5 suit is built. The questions are grounded in existing disclosure requirements: pending proceedings (Item 103), risk factors (Item 105), management's discussion of known trends and uncertainties (Item 303), and loss contingencies that are reasonably possible (ASC 450-20).[130]

---

[126] *Basic*, 485 U.S. at 238.

[127] *See* Evan Hubinger (@EvanHub), X (Sept. 8, 2026) (Anthropic's Alignment Science Lead, writing that "we really do earnestly believe AI could kill all humans" and "I personally think it is >10% within the next decade"); Jacob Coxon (@hilbertspaess), X (Sept. 8, 2026); *see also* Harry Booth, *He Helped Build Powerful AI at OpenAI and Anthropic. Now He's Afraid It Could Kill Us*, Time (Sept. 9, 2026), https://time.com/article/2026/09/09/ai-anthropic-openai-jacob-coxon/; Jon Whittle, *What if AI Doomsaying Is, in Effect, a Cult?*, Austl. Fin. Rev. (Sept. 22, 2026), https://www.afr.com/technology/what-if-ai-doom-saying-is-effectively-a-cult-20260922-p60z0n (defining "p(doom)" as "the probability of AI killing all humans").

[128] *See*, *e.g.*, *Letter from SEC Div. of Corp. Fin. to Carlos Moreira, Chief Exec. Officer*, *SEALSQ Corp.* (Apr. 29, 2024) (directing the issuer to "revise so that your disclosure is consistent with your public statements and promotional materials"); *Letter from SEC Div. of Corp. Fin. to Aptera Motors Corp.* (July 13, 2023) ("we note your public statements regarding your partnership … [p]lease revise to disclose …").

[129] Under the federal securities laws, the investor with the largest financial interest in the relief sought is presumptively appointed lead plaintiff in a securities-fraud class action; public pension funds, often represented by a state attorney general or state fund counsel, frequently fill that role. *See* 15 U.S.C. §§ 77z-1(a)(3), 78u-4(a)(3).

[130] 17 C.F.R. §§ 229.103, 229.105, 229.303; Accounting Standards Codification § 450-20 (Fin. Accounting Standards Bd.).

A. Questions to Ask About Any Frontier AI IPO

Management alignment with shareholders:

- Does the company use dual-class shares? Will founders retain voting control after the IPO?
- Does the company have a traditional 180-day lock-up after the IPO for share sales by insiders? If insiders can sell shares earlier, when can they sell? On what schedule, and how many shares could potentially be sold before the traditional 180-day lock-up period expires?
- Does the company require shareholders to waive class actions?
- Does the company require shareholders to waive jury trials?
- Does the company impose mandatory arbitration on shareholders, rather than letting them file lawsuits in court?
- Does the company impose mandatory arbitration provisions on model users? How does the company view the scope of that arbitration commitment extending beyond users who are in privity of contract with the company?

Financial risks:

- What payments are due within the next 12 months? (This enables comparison of current burn rate with future costs.)
- What payments are due on an annual basis?
- What is the company's cash position as of the most recently ended quarter?
- What share of the company's capital comes from firms that are also its suppliers, customers, or competitors?
- What are the total noncancelable commitments, broken out by party and by year, and what happens to liquidity if demand falls by 5%, 10%, or 20%?
- Which investments depend on milestones, what are those milestones, and has the company met them? What happens if they are not met?
- What compute is financed through leases or special-purpose vehicles that remain off the balance sheet?
- If the company cannot meet its payment obligations, can it renegotiate or restructure its compute contracts, and on what terms?[131]

Market and supplier concentration risks:

- How concentrated are the company's customers?
- How much revenue comes from partners that invest in the company?
- How dependent is the company on any single chip, cloud, or energy supplier, and what contingency plans exist if a supplier defaults or curtails supply?
- Please disclose all equity investments, loans, and other financing provided to the company from cloud service providers. Please identify whether such terms were conditioned on commitments to purchase compute, hardware, or services.

[131] Warren Letter, *supra* note 70.

- Please disclose the amount of revenue the company recognized from cloud service providers' purchases for its own internal use or products, excluding sales to third-party customers
- Please disclose the amount of revenue recognized from third-party sales through its cloud platform including fees, rebates, or other compensation paid to the company on those sales.
- How much of the company's total revenue and year-over-year revenue growth is attributable to customers with no equity investment or financing relationships with the company?

Safety and cybersecurity risks:

- What steps, if any, has the company taken to implement the National Institute of Standards and Technology's (NIST) AI Risk Management Framework?[132]
- What internal estimates exist of the probability and magnitude of the catastrophic risks the company describes?
- What testing supports the company's safety claims, and has that testing has been independently verified?
- Which safety commitments are legally binding, and which can management or the board change on its own?
- Who within the company oversees model testing, and who do they report to?
- What would cause the company to pause development, and who can overrule a decision to pause or limit deployment?
- Provide all internal documents reflecting policies or protocols for determining whether models are safe for deployment.
- How will directors weigh safety obligations against shareholder returns, and how will shareholders learn when the two conflict?
- How does the company determine when to report safety and cybersecurity incidents to investors and relevant governmental authorities?
- Please disclose all safety and cybersecurity incidents. What systems, data, and/or credentials were accessed, obtained, altered, transmitted, or exposed without authorization? Have all affected parties been appropriately notified?
- Were the company's representations to investors and the public about its security practices consistent with the controls actually in place at the time of incidents?
- How were safeguards tested ahead of the incidents?
- Are there incidents the company is aware of but not yet reported publicly?
- What controls has the company adopted to limit or prevent agents from communicating with one another, accessing the internet, or interacting with third-party systems?

[132] NIST, AI Risk Management Framework, https://www.nist.gov/itl/ai-risk-management-framework.

- What steps have been taken to address concerns that sandbox testing environments are not properly isolated or air-gapped?
- What third party tools or systems were used in any safety incidents? Is the company still using those tools or systems? What steps has the company taken to mitigate risks from those third party tools or systems?
- How do contracts between the company and third parties allocate liability? Is the company indemnified? Under what circumstances?
- Identify all instances where the company did not actively monitor AI model tests.
- Please disclose the company's safety protocols and whistleblower policies.
- Will additional funding from an IPO exacerbate safety risks and concerns? Why or why not?
- Has the company reduced spending on safety in the months leading up to the IPO to reduce costs? What forms has that reduction taken, such as reducing headcount or safety research?
- Are there other steps the company has taken to implement or coordinate a "slow down" of development?

Misuse and automation:

- How many misuse incidents has the company detected over the past two years, broken out by category—cyber, biological, chemical, and state-actor?[133] How did the company respond in each case? What has it done to mitigate the risk of future incidents?
- To what extent are AI systems, including the company's own models, used to conduct model research, training, and evaluation, and at which stages does company policy require human review or approval before work proceeds?
- Can those requirements be changed by management alone, and would investors be told if they were?

Legal liability and insurance:

- List all pending legal proceedings, grouped by legal theory—e.g., copyright, product liability, unfair or deceptive acts or practices, defamation, consumer protection, and state statutory claims—and state the reasonably possible loss range for each.
- How would state liability regimes aimed at AI—such as statutes imposing per-violation penalties for nonconsensual imagery or for harms to minors—affect the company's exposure? Provide an estimate for the range of financial exposure. Has government action ever interrupted or altered its operations?
- How do these risks affect insurability, and have any insurers denied or excluded coverage? What exclusions exist?

[133] Cf. Debevoise & Plimpton, *Public Cyber 8-K Tracker* (Sept. 28, 2026) (surveying how issuers disclose the nature, scope, operational impact, and materiality of cybersecurity incidents under Item 1.05).

- Identify all of the ways in which the firm attempts to limit its liability contractually, for example, through warranties or limited liability provisions.
- Quantify the maximum reasonably foreseeable liability from a catastrophic or large-scale harm caused by the firm's models, and state whether current capital and insurance could absorb it. If not, what is the plan—additional reserves, coverage, or contractual indemnities?

B. Questions Specific to Anthropic

Anthropic's S-1 is now on file. But as an initial registration statement it still leaves key figures blank, to be supplied by amendment, so several of the questions below test whether the amended filing quantifies what the registration statement so far describes only in structure, and confirms what reporting on the draft prospectus has already put in numbers.[134] Other questions explore whether there are untrue statements or omissions of material facts that may warrant an SEC action to issue a "stop order" of the initial public offering.

Financing:

- The draft reportedly discloses more than $600 billion in compute commitments—and as much as $750 billion counting deals still under negotiation. What portion is noncancelable and what is the payment schedule?
- What share of Anthropic's revenue and net receivables comes from entities that are at the same time its investors, lenders, or compute suppliers—Amazon, Google, Microsoft, and NVIDIA—and how does the company distinguish genuine third-party demand from spending by its own capital providers?[135]
- Roughly 48% of Anthropic's revenue and 58% of its net receivables now run through the cloud marketplaces of Amazon, Google, and Microsoft. Is this revenue recognized gross or net of the fees remitted to those providers, and are those fees booked as cost of revenue or as sales and marketing? Restate revenue and gross margin treating the revenue-share payments as a cost of revenue. What are the payment and remittance terms, and what recourse does Anthropic have if a platform delays payment or changes them?[136]
- The S-1 describes Anthropic's off-balance-sheet, third-party chip-leasing arrangements only in general terms. Which entities actually hold the leased hardware,

[134] Anthropic S-1, *supra* note 43; *see also Anthropic's IPO Prospectus Shows Sweeping AI Vision, Surging Costs*, Reuters (Sept. 29, 2026), https://www.reuters.com/business/finance/anthropics-ipo-prospectus-shows-sweeping-ai-vision-surging-costs-2026-09-28/.

[135] Anthropic S-1, *supra* note 40 (reporting that no single customer accounted for more than 10% of revenue, with the top five representing approximately 25%).

[136] *Id.* (revenue through cloud-service-provider marketplaces representing approximately 48% of total revenue, and 58% of net accounts receivable, for and as of the six months ended June 30, 2026; platforms remit revenue net of applicable platform and other fees).

and what guarantees, residual-value, or make-whole obligations does Anthropic owe them?[137]

- Anthropic reportedly paid no federal income tax in 2025.[138] In what year does it expect to begin paying federal income tax, and how much did it claim in research-and-development deductions and credits in 2025, broken down by activity—including amounts attributable to model development and to data-center construction and operation?

Governance:

- Anthropic's reported structure gives its founders voting control and lets a "Long-Term Benefit Trust" elect a majority of directors. What voting and information rights will public shareholders have, and when may the company lawfully subordinate their financial interests to those of the Trust?[139]
- Anthropic's lock-up arrangements reportedly let certain insiders sell their shares on the first day of trading and release the remainder on a front-loaded, staggered schedule over roughly two and a half months, rather than the customary 180 days.[140] What number and percentage of shares are eligible to be sold at each tier, and on what dates? Why depart from a single customary lock-up, and do these arrangements heighten the risk of volatility, manipulation, or of insiders exiting ahead of anticipated difficulties?
- How do Anthropic's public statements about the risks of its technology align with its fiduciary duty under Delaware law and its status as a public benefit corporation?[141]

Safety and misuse:

- Anthropic has disclosed five cases of users attempting to misuse Claude for biological research, an unauthorized-access incident involving a pre-release model, and that it "can no longer confidently rule out" dangerous biological uplift. What is the reasonably possible loss and remediation estimate for each incident?[142]
- Anthropic's own CEO has forecast that autonomous AI "swarms" could cause hundreds of billions of dollars in damage within a year. Does the S-1 disclose this risk with the specificity the company's public statements suggest, and what controls and timelines does it commit to?[143]
- Anthropic states in its S-1, "Concerns have been raised both within our organization … that sufficiently advanced AI systems could, if not developed, deployed, and operated with appropriate safeguards, oversight, and security measures, pose

---

[137] *Id.* (describing the third-party financing and chip-leasing arrangements in general terms).
[138] Warren Letter, *supra* note 70.
[139] *See Anthropic's IPO Asks Investors to Back a Company That May Overrule Them*, *supra* note 36.
[140] Warren Letter, *supra* note 70; Anthropic S-1, *supra* note 40, at 307.
[141] *See* Anthropic S-1, *supra* note 43 at 2.
[142] Anthropic, *Detecting and Countering Misuse of AI: September 2026* (Sept. 10, 2026), https://www.anthropic.com/threat-intelligence-report-september-2026; Anthropic, *Activating AI Safety Level 3 Protections* (May 22, 2025), https://www.anthropic.com/news/activating-asl3-protections.
[143] Amodei, *We Must Pace the Frontier*, *supra* note 110.

catastrophic or existential risks to humanity." Has Anthropic assessed the likelihood of that risk, relative to policies and procedures for safeguarding models (which are also omitted from the S-1)? Has Anthropic developed any threat modeling or other analysis to determine the nature of a "catastrophic or existential risk"? What form might that risk take? What specific policies and procedures has Anthropic adopted to prevent all forms of a "catastrophic or existential risk"?

- Describe what steps Anthropic is taking to address its stated inability to provide assurances that "the safeguards or access regimes we adopt will prove adequate, that we will correctly identify which capabilities require restriction, or that constraints we impose will not be circumvented."

Legal, liability, and insurance:

- The FTC has reportedly issued civil investigative demands to Anthropic over safety incidents.[144] Identify all pending government investigations or actions that exist, and describe how could they affect operations.
- How does Anthropic plan to compensate for harms that exceed the capacity of any single insurer or the insurance industry in the aggregate? Has Anthropic been unable to obtain insurance against all potential risks? What types of coverage exclusions has Anthropic been subjected to? How specifically does Anthropic foresee the insurance industry evolving to provide insurance products, and what risk or risks would those policies cover? Is Anthropic able to acquire insurance for a "catastrophic or existential risk"?
- What is the basis for Anthropic's stated uncertainty about the enforceability of its warranty disclaimers and limitations of liability in its terms of service? Identify all lawsuits and pending legislation relevant to this assessment.

Conclusion

A frontier AI IPO of this scale will place its risks—financial, legal, and catastrophic—into the retirement accounts of ordinary Americans within days of listing, largely regardless of their choosing. The SEC's disclosure review, backed by its authority over the timing and effectiveness of the offering, is the point at which those risks can still be surfaced and priced. The questions above are designed to help the Commission, and the fiduciaries who will be forced buyers, obtain the disclosure that the securities laws require before the public is exposed. Fund managers, fiduciaries, and other investors should watch carefully for the figures the S-1 still leaves blank, and if they are not filled, press for disclosure through the comment process.[145]

---

[144] *See FTC Issues Subpoenas to Anthropic and OpenAI over Safety Incidents*, CBS News (Oct. 15, 2026), https://www.cbsnews.com/news/ftc-subpoenas-anthropic-openai-consumer-protection/.

[145] *Id.* (also omitting the amounts due before 2030 and within the next year, providing no counterparty breakdown, and describing the off-balance-sheet chip-leasing arrangements only in general terms).