Who Holds Correlated AI Losses?

Concentration and Risk Allocation in the AI Supply Chain

Ilker Aslan[1]

November 2026

Abstract

Firms that use artificial intelligence buy their models from very few providers, so choices of provider that each firm makes on its own tie the liability risk of many firms to a few common shocks. I study how this structure of the supplying industry decides who bears the losses. However many policies an insurer writes, the AI component of its book carries the systematic variance of only $1/H$ independent risks, where H is the Herfindahl index of provider shares. Diversified capital could carry this accumulation cheaply if slices of it could be traded, but liability wordings do not identify the provider behind a claim, and insurers have begun to exclude AI as such. An excluded loss stays with the firm that deployed the model and, if the firm defaults, passes to its bank. I then measure what supervisors do not observe. Of 313 firms whose annual reports name a provider, 89 say that their products run on its models, 73 of them on OpenAI's; 22 of these firms and four AI-infrastructure borrowers owe USD 78 billion in syndicated loans, held to a large extent by the same banks. When a provider has an incident, its clients' shares fall by 1.4 percentage points over two days relative to the clients of its competitors. Their loans carry no spread premium and their bonds do not move. Contracts filed with the SEC show how the factor is allocated: providers are protected against their suppliers' failure, deployers report neither recourse nor cover, and the suppliers' lenders hold the provider's credit under covenants tied to its rating. For 104 European banks, a defect in the largest provider's model that harms a tenth of its clients would cost the median bank 0.3% of its CET1 capital, or 0.7% if its AI-using borrowers all rely on that provider. Equity prices register the provider factor; loan and bond prices, credit files and capital rules do not. I propose recording provider dependence in credit and underwriting files and running one provider-failure scenario across insurance and banking supervision.

JEL classification: G14, G21, G22, G28, K13, L86, O33.

Keywords: artificial intelligence, AI adoption, foundation models, market concentration, accumulation risk, liability insurance, exclusions, syndicated loans, event study, financial stability.

[1]Lecturer, Kalaidos University of Applied Sciences, Zürich. E-mail: ilker.aslan@alumni.unifr.ch. ORCID: 0000-0001-9241-9420. Code and data: doi:10.5281/zenodo.23733783. All errors are mine.

1 Introduction

A small number of firms supply the models on which the rest of the economy is building its use of artificial intelligence: at the end of 2025, three providers took 88% of enterprise spending on model interfaces (Tully et al. 2025b), while the share of EU enterprises with 250 or more persons employed that use AI generating written or spoken language rose from 8% in 2023 to 32% in 2025 (Eurostat 2026a). Firms route customer service, credit and hiring decisions, software development and the drafting of contracts through these models. Each firm chooses its provider for its own reasons, but a defect in a model reaches all of its clients at once, so together these choices tie the liability of many firms to a few common shocks. This paper asks how the structure of the supplying industry decides who bears those losses: the providers, the firms that deploy the models, their insurers, or the banks that lend along the supply chain.

The paper establishes three findings, which Figure 1 places along that chain. First, markets see the provider factor. Over the day of a provider incident and the next, the shares of firms whose products run on the affected provider's models fall by 1.42 percentage points relative to those of firms that rely on other providers ($t = -3.6$ across 38 incidents), and none of 1,000 sets of random dates produces a difference as large. Second, credit does not. The spreads of syndicated loans to the same firms are no higher than those of comparable loans, with a confidence interval that excludes a premium above about 43 basis points; their bonds do not move at incidents and their yields carry no premium; and the internal ratings-based formula gives the same capital to a lender whose borrowers all rely on one provider as to one whose borrowers are spread across providers. Third, contracts move the factor away from the providers. Providers' terms cap their liability to business customers and indemnify them only against intellectual-property claims; of the 114 firms that disclose dependence in their annual reports, 65 name it as a risk, but none reports a service-level commitment or insurance that covers a provider's failure. The agreements that suppliers of computing capacity file give the providers step-in rights, service credits and supplier indemnities, while the suppliers' lenders take the provider's credit directly: a CoreWeave facility for OpenAI capacity traps cash if OpenAI stops paying, and CoreWeave's guarantee falls away once OpenAI is rated investment grade, in loans that credit files record as exposures to CoreWeave.

A model explains why the factor ends where it does. I model firms that adopt AI from one of a few providers, the insurers that cover their liability and the banks that lend to them. A failure at a provider is a common shock to its clients, so no number of policies or loans diversifies it away: the AI component of a portfolio retains a systematic variance set by the Herfindahl index of provider shares (Lemma 1), and at end-2025 shares a book of any size diversifies its AI risk as though the market were served by about 3.6 equal providers. Measured on its own, the accumulation needs capital close to the largest provider's share of a full-market event, and cover then survives only if insureds' surplus from it is at least 18 to 22% of expected losses (Proposition 1). Inside a large balance sheet with unrelated risks it needs little: if slices of it could be traded, its price would be roughly proportional to the Herfindahl index divided by the diversified capacity that holds it (Proposition 2), and across the US property-casualty sector it would cost less than a hundredth of what would make insureds forgo cover. What keeps the accumulation with the carriers that write it are frictions such as the price of catastrophe reinsurance (Froot 2001) and wordings that do not identify

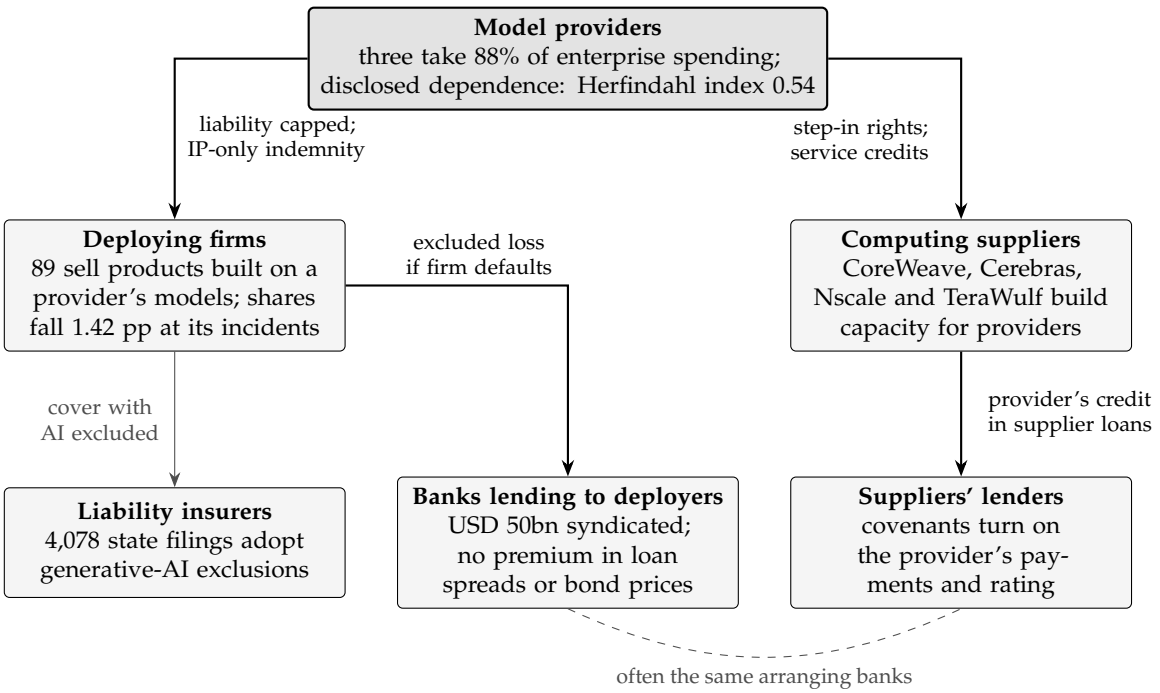


Figure 1. Where a provider failure lands

Notes: Solid arrows show the direction in which a loss from a provider failure travels under the contracts and cover in force; the labels on them summarise the terms that send it there. Figures from Sections 2, 5 and 6.

the provider behind a claim. Carriers whose books it dominates exclude it, and because no firm, insurer or bank internalises what correlated defaults cost the credit supply of firms that use no AI, exclusion can be privately chosen and socially inefficient (Proposition 1).

Insurers have begun to exclude. Standard exclusions for generative AI, published for the US commercial general liability form in July 2025, had been adopted in 4,078 state filings a year later (Rosanes 2026), most of them removing AI as such rather than limiting exposure per provider (Section 2.4). A loss that is excluded stays with the firm that deployed the model and, if the firm defaults, passes to its lender (Corollary 1). For 104 European banks, in a year in which the largest provider fails with none of the AI-related liability of borrowers insured, the median bank would lose 0.3% of its CET1 capital if its borrowers mirror the market and 0.7% if all its AI-using borrowers rely on the failing provider, and the most exposed bank 2.2% (Section 5). What drives these losses is the ratio of a bank's corporate loans to its capital, which supervisors observe, and the share of those loans that go to clients of the failing provider, which they do not. Banks hold the factor a second time as lenders to AI infrastructure whose contracted revenue comes from the same providers, and at a loss rate of an eighth on that lending the asset channel matches the whole liability channel once it reaches 1.7% of a bank's corporate book where its AI-using borrowers rely on the failing provider. The scenario is severe and its severity is American, scaled on US liability premium; at claims per harm half that severe, the bank figures fall to a fifth.

The data behind these findings come from public filings, commercial loan and bond data and share prices (Section 6). Of 313 firms whose annual reports from 2023 to 2026 name a model provider, 89 state that products they sell run on its models; 73 of them name OpenAI, so that disclosed dependence, with a Herfindahl index of 0.54, is more concentrated

than spending. Twenty-two of the firms that disclose dependence, together with four AI-infrastructure borrowers, owe USD 78 billion in syndicated loans, and 19 of the 30 lenders to the infrastructure firms also lend to firms that disclose dependence. Firms diversify slowly: between successive annual reports ten add a provider, every one keeps the provider it had, and firms that name several providers fall as much at their provider's incidents as those that name only one.

Credit practice and supervision are not set up to see the factor. Work on AI and financial stability concerns the AI that banks and insurers use themselves (Financial Stability Board 2024; European Systemic Risk Board 2025; Bank of England 2025; Buch 2026), the oversight of critical technology providers under the Digital Operational Resilience Act leaves out the dependence of financial institutions' borrowers and policyholders on the same providers (European Supervisory Authorities 2025), and credit files record a borrower's industry, size and location but not the models its operations depend on. In the ECB's bank lending survey of October 2026, only a small number of banks said that firms' adoption of AI features in their assessments of borrowers (European Central Bank 2026c). I propose that credit and underwriting files record which model a firm's core processes depend on, and that insurance and banking supervisors run a single provider-failure scenario under the contract wordings in force, which would show in which sector the loss lands (Section 7). The revised Product Liability Directive, which applies from December 2026, moves part of the loss to providers only as far as their capacity in the state in which they fail allows (Appendix G).

Related literature. The paper applies standard tools, among them diversity scores (Cifuentes and O'Connor 1996), coherent risk measures (Artzner et al. 1999), the Euler allocation of capital (Tasche 2008; Kalkbrener 2005), the single-factor model of credit risk and the judgment-proof problem, to a common factor created by the structure of the AI supply industry.[2] Work on AI and financial stability has emphasised the uniformity, opacity and procyclicality of the models that financial institutions use and concentration among their vendors (Gensler and Bailey 2020; Daníelsson et al. 2022; Financial Stability Board 2024; Aldasoro et al. 2024; European Systemic Risk Board 2025; Leytes 2026), following work on cloud dependencies (Financial Stability Board 2019), and shared models homogenise outcomes among the firms that use them (Kleinberg and Raghavan 2021; Bommasani et al. 2022; Toups et al. 2023). On insurance, Leung et al. (2026) single out foundation-model concentration as the clearest new limit on the insurability of AI, Trout et al. (2026), Zhu (2026) and Cloud Security Alliance (2026) propose aggregates and dependency mapping by provider, and Lopez and Nkameni (2026) calibrate cloud-outage scenarios for cyber insurers. Uninsured catastrophe losses reach bank credit (Garmaise and Moskowitz 2009; Cortés and Strahan 2017; Koetter et al. 2020), banks' operational losses are correlated across institutions (Berger et al. 2022), shocks propagate through common suppliers (Barrot and Sauvagnat 2016; Carvalho et al. 2021), and climate stress tests trace losses through borrowers to their lenders (Battiston et al. 2017). Cyber risk, whose losses also correlate through common technology (Biener et al. 2015; Eling and Wirfs 2019; Welburn and Strong 2022; European Systemic Risk Board 2020; Duffie and Younger 2019; Kashyap and Wetherilt 2019; Aldasoro et al. 2022), supplies the empirical designs

[2] In Aslan (2026a) I study catastrophic AI risks that may be uninsurable at any price; the risk here can be priced but is correlated across insureds.

that dependence data allow (Florackis et al. 2023; Jamilov et al. 2021; Eisfeldt et al. 2023; Crosignani et al. 2023; Eisenbach et al. 2022). The paper also builds on work on insurability and catastrophe capital (Berliner 1982; Jaffee and Russell 1997; Froot 2001; Zanjani 2002; Ibragimov et al. 2009, 2011), concentration in credit portfolios (Pykhtin 2004; Düllmann and Masschelein 2007; Gordy and Lütkebohmert 2013) and judgment-proof injurers and liability for AI (Shavell 1986; Pitchford 1995; Boyer and Laffont 1997; Lior 2022; Hacker 2023; Buiten et al. 2023). Relative to work on sector concentration, the sector here is defined by the provider a borrower depends on rather than by its industry; relative to climate stress tests, the shock is a failure at one provider that reaches all its clients in the same year, across industries and regions; and relative to work on AI adoption, the paper follows the risk that adoption creates through the contracts and credit of the supply chain rather than its effect on productivity.

Section 2 presents the institutional evidence and Section 3 the model. Section 4 quantifies the cost of carrying the accumulation in insurance and Section 5 the exposure of banks. Section 6 measures dependence, lenders' holdings, the response of share, loan and bond prices, and the terms of the contracts with providers. Section 7 discusses policy and Section 8 concludes.

2 Provider dependence and who has borne its losses

2.1 A concentrated supply of models

Three providers took 88% of enterprise spending on model interfaces at the end of 2025 (Anthropic 40%, OpenAI 27% and Google 21%), and in 2023 OpenAI alone took half, according to Menlo Ventures' surveys of enterprise buyers (Tully et al. 2025a,b). Menlo Ventures is one of Anthropic's largest investors (Temkin 2024), and other sources disagree about which provider leads, but not about concentration: the largest provider's share is about 56% in a survey of large firms in early 2026 (Wang et al. 2026) and 61% in transaction data for the week to 23 October 2026 (Ramp 2026), and the Herfindahl index lies between 0.21 and 0.37 across the survey snapshots of Figure 2a and the splits of their residual shares, and at 0.49 in the transaction data (Tables 10 and 14 in Appendix C). The share of EU enterprises with ten or more persons employed that use AI generating written or spoken language, the closest official proxy for reliance on foundation models, rose from 2% in 2023 to 9% in 2025, and among enterprises with 250 or more from 8 to 32% (Figure 2b). Of the annual reports filed on Form 10-K in 2026, 60% mention artificial intelligence and 454 refer to third-party AI, almost six times as many as in 2024, but only 41 of those 454 name OpenAI, ChatGPT or Anthropic (Table 11 in Appendix C), a count that omits model names such as Claude and Gemini. Outside observers cannot tell which firms depend on which provider.

Where firms do name a provider, their lenders can be traced. Of 93 firms that name OpenAI or Anthropic as a dependency in filings from January 2025 to October 2026, 31 name the administrative agent of a credit facility, and one bank is agent to 11 of them (Table 12 in Appendix C; details in Appendix H.1). Section 6 follows a larger set of firms, identified from annual reports, to every member of their lending syndicates and the share each holds.

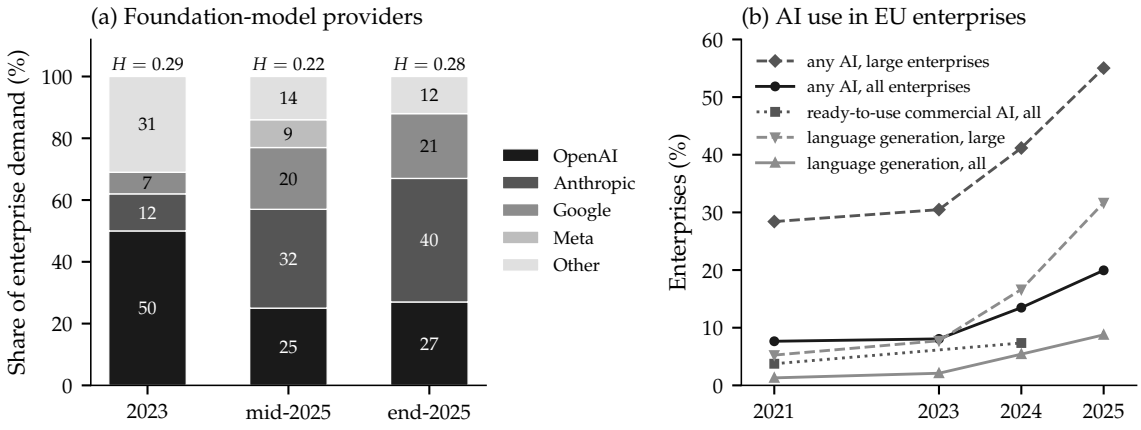


Figure 2. Provider shares and the use of AI by EU enterprises

Notes: Panel (a): shares of enterprise spending on foundation-model APIs in 2023 and at the end of 2025, and shares of enterprise production usage in mid-2025, from Menlo Ventures' surveys of enterprise buyers (Tully et al. 2025a,b). Meta is included in "Other" where it is not reported separately. H is the Herfindahl index with the residual share divided equally among four providers. Panel (b): share of enterprises in the EU27 with ten or more persons employed ("all") or 250 or more ("large") that use at least one AI technology, that use AI generating written or spoken language, and that use ready-to-use commercial AI systems, a question asked in 2021 and 2024 only (Eurostat 2026a). There is no observation for 2022.

2.2 How often providers fail

The events closest to a liability-relevant failure are changes in a model's outputs that reached the customers of a provider's interface. At Anthropic, OpenAI, Google and Meta, the four providers with the largest shares in the surveys, two occurred from January 2023 to October 2026, in 13.0 provider-years if Meta is counted from April 2025, when it began to offer its own interface, and in 15.3 if it is counted throughout: a tendency of one model to leave tasks incomplete, addressed in an update in January 2024 (OpenAI 2024), and three infrastructure bugs that degraded one provider's responses in August and September 2025, took between eight and 44 days to fix and reached about 30% of the users of one coding product (Anthropic 2025a). Neither is documented to have produced a liability claim from a business customer, so a liability-relevant failure remains a scenario and its probability a parameter (Appendix C.1).

The two belong to 55 events at eight providers in which one change or failure reached many clients at once (Table 21 in Appendix D). Sixteen were changes in a model's behaviour that the provider rolled back or acknowledged as defects, half of them reversed within about a week and six lasting from two weeks to six months. Apart from the two above, they were confined to consumer applications, third-party deployments or other products, or the provider's account does not say whether they reached its interface. Broad outages lasted a median of five to six hours, and events can also be regulatory: on 12 June 2026, access to two frontier models was suspended for all customers to comply with a US export-control directive (Anthropic 2026b,a). The only event with published loss estimates, an outage of Amazon Web Services in October 2025 (Amazon Web Services 2025), cost an estimated USD 38 to 581 million of insured losses, all from downtime (CyberCube 2025; Insurance Edge 2025). When OpenAI's governance crisis of November 2023 put its continuity in doubt for three trading days, the shares of firms that had disclosed using its models did not fall relative to those of firms that mention it only as an example, but with 21 and 8 such firms a discount below about

3.5 percentage points could not be detected (Appendix E). Across 38 dated incidents from 2023 to 2026, however, the shares of firms whose products run on the affected provider's models fell relative to those of firms that rely on other providers, by 1.4 percentage points over the incident day and the next (Section 6.3).

2.3 Where losses have landed

The best-documented failure of a widely used vendor, a faulty update to CrowdStrike's Falcon sensor that crashed about 8.5 million Windows computers on 19 July 2024, left most of its losses with the firms it hit (Weston 2024; Parametrix 2024): 80 to 90% of the USD 5.4 billion direct loss of the affected US Fortune 500 companies, excluding Microsoft, was uninsured, and the vendor, whose terms capped its liability at the fees for the relevant offering, conceded about 1.5% of that loss (Table 8 in Appendix A). The terms of foundation-model providers leave deployers as little recourse. Every provider whose terms could be read in full caps its liability to business customers at about a year of fees or excludes consequential loss, indemnifies them only against intellectual-property claims, and in most cases is itself indemnified by the customer against third-party claims arising from the customer's applications, inputs or outputs (Table 9 in Appendix A). Early cases have reached deployers, vendors and developers alike: a Canadian tribunal held an airline liable for its chatbot's misleading advice (Moffatt v. Air Canada 2024), a US federal court let age-discrimination claims against the vendor of a widely used screening tool proceed as a nationwide collective (Mobley v. Workday 2024, 2025), and a developer of foundation models agreed to pay USD 1.5 billion over the books used to train its models (Bartz v. Anthropic 2026), a claim about training data.

2.4 Insurers' response

When generative AI spread, most liability policies in force said nothing about it, as cyber losses had sat inside property and liability policies never priced for them until the Prudential Regulation Authority and Lloyd's required insurers to identify such non-affirmative exposure from 2017 (Prudential Regulation Authority 2017; Lloyd's 2019, 2020). The Insurance Services Office published optional generative-AI exclusions for the commercial general liability form (CG 40 47 and CG 40 48) and for products and completed operations (CG 35 08) in July 2025; adoptions were filed from August, peaked in February 2026 and reached 4,078 in 49 states and the District of Columbia by the end of July 2026, and seven insurers, among them large multi-line groups, made more than 60% of the filings of early 2026 (Rosanes 2026; Independent Insurance Agents & Brokers of America 2025; Germond 2026). In management and professional liability, carriers write exclusions on their own forms, the broadest barring claims arising from "any actual or alleged use, deployment or development" of AI by any person (Bloomberg Law 2026; Maschke et al. 2026).

Appendix D codes 54 actions by insurers, reinsurers, brokers and form providers from January 2023 to October 2026 (Table 22 in Appendix D): 13 restrict cover, 29 offer affirmative cover and 12 are public statements about accumulation or capacity. Table 23 in Appendix D compares them with what three explanations of the restrictions predict. Nine of the 13 exclude AI or generative AI as such, the wording that carriers removing silent cover would write and also one that a carrier unable to measure its exposure to each provider would choose.

One restriction excludes incidents that cause widespread losses at the same time (Harrington 2026), one draft would cap specified AI losses at about a tenth of the policy limit, and two cap particular perils. Six of the 13 come from specialty carriers, the writers on which the capital mechanism of Appendix C.3 bears, while the seven insurers behind most of the standard-form filings include large multi-line groups, on which it does not. Affirmative liability cover for firms that use AI comes mostly from managing agents and coverholders at Lloyd's, with stated limits of USD 9.25 million to USD 25 million or more per insured, and no document sets an aggregate across insureds for the losses from one provider's failure (Armilla AI 2026; Table 25).

Some market participants tie AI cover to accumulation explicitly. Beazley's chief underwriting officer said in 2023 that the firm would set its appetite for AI from a disaster scenario and could exclude exposure beyond it (Wilkinson 2023), a broker contrasted a USD 400 million loss to one company with a failure that triggers 10,000 losses at once (Loizos 2025), and reinsurance brokers trace the systemic dimension of AI risk to common foundation models (Lockton Re and Armilla AI 2026; Gallagher Re 2026). The absence of any aggregate per provider event, which the accumulation explanation would lead one to expect, counts against it, although the difficulty of attributing claims to a provider could also keep such limits out of the market (Section 3.5); which of the two applies can be tested once dependence is recorded (Appendix H.6). The developers of frontier models have struggled to buy cover: one reportedly arranged up to USD 300 million, a figure another person familiar with the policy put much lower, and two were reported to be weighing the use of investor funds to settle claims (Business Insurance 2025).

2.5 Prudential rules and credit practice

Prudential rules do not record the dependence of policyholders and borrowers on these providers. Neither the Solvency II standard formula for general liability nor its man-made catastrophe scenario for liability depends on the technology that insureds share, bank capital under the internal ratings-based approach rests on a single systematic factor (European Commission 2015; Basel Committee on Banking Supervision 2005), and the oversight of critical technology providers under the Digital Operational Resilience Act concerns the services supplied to financial institutions themselves (Financial Stability Board 2023; European Supervisory Authorities 2025). Supervisors' concern about concentration has the same focus: the Chair of the ECB's Supervisory Board has warned that the financial entities adopting AI rely on just a few technology suppliers, and that AI-related risks could become systemic if failures in widely used AI tools affected many firms at once (Buch 2026). Banks' credit assessments do not yet fill the gap. In the ECB's bank lending survey of October 2026, most of the 159 banks had not introduced specific criteria for assessing lending that finances firms' adoption or development of AI, and only a small number said that the adoption of AI and its risks increasingly feature in their qualitative assessments of borrowers (European Central Bank 2026c). The revised Product Liability Directive, which applies to products placed on the market after 9 December 2026, adds strict liability of model developers towards natural persons (European Parliament and Council 2024). Appendix A gives the detail.

3 A model of correlated AI losses

3.1 Environment

There is a unit mass of firms, each of which buys liability insurance. A fraction $a \in (0,1]$ of firms use AI in their operations. Each adopter relies on one of K providers of foundation models; provider k serves a share σ_k of adopters, with $\sum_{k=1}^{K} \sigma_k = 1$. I measure concentration by the Herfindahl index $H(\sigma) = \sum_k \sigma_k^2$, whose inverse is the number of equal-sized providers that would produce the same index.

Adopters face idiosyncratic and common AI-related harms. With probability π, a firm's own use of AI harms a third party (a misconfigured agent, an unchecked output, an inappropriate use case), independently across firms. Common harms originate at the providers: in each period provider k suffers a failure with probability p_k, and the failure indicators $F_k \in \{0,1\}$ are independent across providers, an assumption relaxed in Appendix F. A failure is any event at the provider that changes the behaviour of the model for all its clients, such as a defective update, a regression in safeguards or a compromised component; for liability the events that matter are changes in what the model produces, and the calibration of p rests on those alone. Conditional on $F_k = 1$, each client of provider k suffers a harm with probability q, independently across clients. Every harm generates a liability loss X, drawn independently from a distribution with mean μ and coefficient of variation δ.

The aggregate state is summarised by the share of adopters whose provider fails,

$$T = \sum_{k=1}^{K} \sigma_k F_k \in [0,1], \tag{1}$$

with mean $\bar{p} \equiv \mathrm{E}[T] = \sum_k \sigma_k p_k$. Aggregate AI losses depend on the state only through T, whose distribution is set by the failure probabilities and by market structure.

An insurer writes n policies on adopters and its book mirrors the market, so that $n_k = \sigma_k n$ policies are written on clients of provider k. A policy covers a fraction $c \in [0,1]$ of AI-related liability losses: $c = 1$ is full cover, whether affirmative or silent, and $c = 0$ is a complete exclusion. The insured AI loss on policy i is $Y_i = c\,(X_i I_i + X_i' J_i)$, where I_i and J_i indicate an idiosyncratic and a common harm, and X_i, X_i' are independent draws of the severity. The average loss per policy is $\bar{Y}_n = n^{-1} \sum_{i=1}^{n} Y_i$.

3.2 Diversification stops at the provider level

Lemma 1 (Diversification ceiling). *(i) As $n \to \infty$, $\bar{Y}_n \to c\Lambda$ almost surely, where $\Lambda = \mu(\pi + qT)$.*

(ii) For finite n,

$$\mathrm{Var}(\bar{Y}_n) = c^2\mu^2 q^2 \sum_k \sigma_k^2 p_k(1-p_k) + \frac{c^2 V}{n}, \tag{2}$$

where $V = \sum_k \sigma_k \,\mathrm{E}\left[\mathrm{Var}(Y_i/c \mid F_k)\right]$, with i a client of provider k, is the average conditional variance of a policy.

(iii) If $p_k = p$ for all k, the systematic variance equals $c^2\mu^2 q^2 p(1-p)H(\sigma)$, as if the market were served by $1/H(\sigma)$ equal providers, however large the book. The effective number of independent

policies, $n_{\text{eff}}(n) = \text{Var}(Y_1)/\,\text{Var}(\bar{Y}_n)$*, increases in* n *towards the ceiling*

$$\bar{n}_{\text{eff}} = \frac{(1+\delta^2)(\pi + qp) - \pi^2 - q^2p^2}{q^2 p(1-p)\, H(\sigma)}. \tag{3}$$

Proofs are in Appendix B. The calculation is the one behind the diversity scores that rating agencies assign to pools of correlated loans (Cifuentes and O'Connor 1996), applied to a common factor set by industrial organisation, and it applies in the same way to a bank's loans to adopters, whose defaults load on T. An individual insurer can do better by over-weighting clients of small providers, since with provider shares $\tilde{\sigma}_k$ in its own book its systematic variance is proportional to $\sum_k \tilde{\sigma}_k^2$, but the books of all insurers add up to the market, so the size-weighted average of these indices is at least H by Jensen's inequality. The equivalence with $1/H$ equal providers holds for the variance only, because tail measures depend on the largest share.

3.3 Capital for the accumulation

In the large-book limit, the capital an insurer needs for the AI component of its book depends only on the distribution of T. With a risk measure ρ applied to annual losses, stand-alone capital per policy is

$$C_\rho = c\,\mu q\,[\rho(T) - \text{E}[T]]^+, \tag{4}$$

because a measure that is translation-equivariant and positively homogeneous is unaffected by the idiosyncratic component, which converges to its mean, and scales with $c\mu q$. I use expected shortfall, $\text{ES}_\alpha(T) = (1-\alpha)^{-1}\int_\alpha^1 \text{VaR}_u(T)\,du$ with $\text{VaR}_u(T) = \inf\{t : \Pr(T \le t) \ge u\}$, the measure of the Swiss Solvency Test with $\alpha = 99\%$ (Swiss Federal Council 2024, Art. 22). Under expected shortfall and with equal failure probabilities, capital does not fall when market share moves from smaller to larger providers, and for any failure probabilities the expected shortfall of T is at least the largest provider's share times the smaller of one and the ratio of its failure probability to the tail probability (Lemma 2 in Appendix B). Value-at-risk, the measure of Solvency II with $\alpha = 99.5\%$ (European Parliament and Council 2009, Art. 101(3)), assigns the accumulation no capital when the probability that any provider fails is below the tail probability, as silent cover implicitly does, and it can rank a market served by one provider as safer than the same market served by several (Lemma 3), the textbook failure of value-at-risk for discrete losses whose probability sits near the confidence level (Artzner et al. 1999; Daníelsson et al. 2013).

3.4 When cover is supplied

Consider the choice between full cover ($c = 1$) and exclusion ($c = 0$) of AI-related losses, idiosyncratic and common alike, which is the choice that an AI exclusion makes. Firms value cover at $(1+v)$ per unit of expected covered loss, where the surplus $v > 0$ stands for the deadweight costs of default and financial distress, the risk aversion of owners and managers, and contractual requirements to carry cover (Mayers and Smith 1982; Froot et al. 1993); the costs of distress and default alone have been estimated at a tenth to a fifth of firm value (Andrade and Kaplan 1998; Davydenko et al. 2012), a different unit from v, which is measured per unit of expected covered loss. There is no estimate of v itself, and Appendix C.1 uses illustrative values. The insurer charges the expected loss, an expense loading η, and a cost of

capital κ on the capital it attributes to the accumulation (Froot and Stein 1998). I start from stand-alone capital (4), which is how insurers that set appetites scenario by scenario assess catastrophe and cyber exposures, as in the realistic disaster scenarios of Lloyd's, one of which is a cascading failure of cloud services (Prudential Regulation Authority 2017; Lloyd's 2026). Writing $h = \pi + q\bar{p}$ for the expected number of AI harms per adopter, cover is traded if and only if $(1+\eta)\,\mu h + \kappa\,\mu q\,[\rho(T) - \bar{p}]^+ \le (1+v)\,\mu h$. Dividing by $\kappa\mu q$ and writing $r = \pi/q$ gives a condition on an index that depends on the distribution of T and, through r, on the relative weight of idiosyncratic harms:[3]

$$\Xi(\sigma, p) \equiv \frac{[\rho(T) - \bar{p}]^+}{r + \bar{p}} \;\le\; \frac{v - \eta}{\kappa}. \tag{5}$$

The numerator is the capital that the accumulation requires, in units of a full-market event; the denominator is the expected loss that the insurer can charge for. Treating v as constant ignores that cover is worth more to a firm in states in which credit also tightens, which would raise the private value of covering the accumulation.

The index does not depend on how large the accumulation is, because capital and expected loss both scale with adoption and severity. Insurers' appetites are set relative to their own capital, however. Write P for the premium of a book, a_w for the share of that premium written on adopters and $z = q\mu/P_i$ for the expected AI-related liability loss of a client of a failing provider in units of its own premium P_i, so that $a_w z$ is the loss of a full-market event as a share of a year's premium. An insurer that accepts a single scenario only while the capital it attributes to the scenario, or the loss of the scenario itself, stays below a share ϵ of its equity E also requires

$$a_w z\, C_s \le \epsilon\, E/P, \qquad C_s = [\rho(T) - \bar{p}]^+ \text{ or } C_s = \max_k \sigma_k, \tag{6}$$

where the second form is the loss from a failure of the largest provider, the scenario that a realistic-disaster exercise would select.

An exclusion moves the loss without removing it: a firm with loss-absorbing capacity B that bears an uninsured harm with loss X defaults when $X > B$, which happens with probability $\Psi_0 = \Pr(X > B)$, and its default causes a loss to its bank. Banks' credit losses under exclusion, $D_0(T)$, given by (10) below at $c = 0$ with b the lending per adopter and ℓ the loss given default, are perfectly correlated with T; under full cover they are zero. Expected losses on loans are priced into lending rates and borne by firms ex ante, so they are part of v. The systemic cost of correlated defaults is not priced: as losses in a state erode banks' capital, banks cut lending to all borrowers, including firms that use no AI, and most sharply once losses exceed the capital they hold above their requirements. I represent this cost, which no individual firm, insurer or bank internalises, by a convex, non-decreasing function $\Gamma(D)$ with $\Gamma(0) = 0$, in the spirit of the externalities that motivate macroprudential policy (Lorenzoni 2008; Bianchi 2011).

Proposition 1 (Exclusion)**.** *(i) Cover is supplied and bought if and only if $\Xi(\sigma, p) \le (v - \eta)/\kappa$ and, for an insurer with a scenario appetite, condition* (6) *holds.*

(ii) If $p_k = p$ for all k and capital is set by expected shortfall, Ξ and the left-hand side of (6) *are*

[3] Condition (5) has the form of the insurability frontier in Aslan (2026d), where capital is set under ambiguity about the loss distribution.

Schur-convex in σ, and the latter rises in proportion to $a_w z$. A more concentrated provider market, or wider or more intense use of AI, is weakly closer to exclusion.

(iii) Cover is socially efficient if and only if $\Xi(\sigma, p) \leq (v - \eta)/\kappa + W(\sigma)$, where

$$W(\sigma) = \frac{\mathrm{E}\left[\Gamma\left(D_0(T)\right)\right]}{\kappa\, a\mu q\,(r + \bar{p})} \geq 0. \tag{7}$$

Exclusion is privately chosen but socially inefficient in particular when $(v - \eta)/\kappa < \Xi \leq (v - \eta)/\kappa + W$, and whenever (6) alone rules out cover while $\Xi \leq (v - \eta)/\kappa + W$. If $p_k = p$ for all k, the wedge W is Schur-convex in σ.

With equal failure probabilities, neither the private cost of covering AI nor the social cost of excluding it falls when concentration rises. If Γ is linear, with slope γ, the wedge $W = \gamma \ell b \Psi_0/(\kappa\mu)$ does not depend on shares; it responds to concentration only when the cost of bank losses is convex, for instance because losses beyond banks' voluntary buffers force cuts in lending. Whether the wedge is large is an empirical question, to which Section 5 returns. Between full cover and exclusion lies a cap on what the insurer pays when any one provider fails, which can bring the accumulation within an insurer's appetite while the insured firms keep its tail (Appendix F).

3.5 Who can hold the accumulation

Stand-alone capital is the limit of integrated capital as the accumulation grows relative to the rest of the book that holds it. Consider a carrier that holds an amount $\theta \geq 0$ of the market's AI accumulation, measured in units of the whole accumulation, alongside an independent book of other risks whose annual loss R, measured like T in units of a full-market event, has a density and a finite mean. With ρ expected shortfall at level α, the carrier's Euler capital and its incremental capital per unit of the accumulation it holds are

$$e(\theta) = \mathrm{E}\big[T \mid \theta T + R \geq \mathrm{VaR}_\alpha(\theta T + R)\big] - \mathrm{E}[T], \qquad \bar{e}(\theta) = \frac{\rho(\theta T + R) - \rho(R)}{\theta} - \mathrm{E}[T]. \tag{8}$$

The Euler capital is the marginal capital of the accumulation (Tasche 2008), and incremental capital is its average over the amounts held, $\bar{e}(\theta) = \theta^{-1} \int_0^\theta e(u)\, du$. Both rise with θ, from zero for a small slice to the stand-alone capital $\rho(T) - \mathrm{E}[T]$ for a holding that dominates the book, and with equal failure probabilities incremental capital does not fall when concentration rises (Proposition 3 in Appendix B). Because R enters only through its scale when its shape is fixed, both depend on θ relative to the standard deviation of R. A large reinsurer that holds a slice of the accumulation alongside unrelated perils needs almost no capital for it, while a carrier whose book is dominated by AI-exposed insureds needs close to stand-alone capital under any internal model.

Sharing the accumulation. Let carriers $j = 1, \ldots, M$ hold books $R_j = s_j U_j$, where the U_j are independent of each other and of T, identically distributed with a density and a finite mean, and $s_j > 0$ is the scale of carrier j's other risks in units of a full-market event, and write $S = \sum_j s_j$. Write e_1 and $\bar{e}_1$ for the capitals in (8) of a carrier with $s = 1$, $\theta_1^* = \sup\{\theta : \bar{e}_1(\theta) \leq (r + \bar{p})(v - \eta)/\kappa\}$ for the largest amount of its own accumulation that such a carrier

covers at the private threshold, and $\bar{m} = (v - \eta)(r + \bar{p})$ for the largest margin per unit of the accumulation that insureds accept for full cover. Slices of the accumulation can be traded among carriers, through quota shares, event covers or reinsurance, at a margin m per unit above expected loss, and each carrier takes m as given.

Proposition 2 (Who holds the accumulation)**.** *(i)* Allocation. *Every carrier holds the same multiple of the scale of its other risks: carrier j holds $\theta_j(m) = s_j\, g(m/\kappa)$, with $g(t) = \sup\{u \geq 0 : e_1(u) \leq t\}$.*

(ii) Price. *Cover is supplied and bought if and only if $\kappa e_1(1/S) \leq \bar{m}$. The accumulation then trades at the margin $m^* = \kappa e_1(1/S)$, and carrier j holds s_j/S of it.*

(iii) Price of concentration. *If the U_j are standard normal and $p_k = p$ for all k,*

$$m^* = \kappa\, \frac{\phi\left(\Phi^{-1}(\alpha)\right)}{1-\alpha}\, \frac{p(1-p)\, H(\sigma)}{S} + o(1/S) \qquad \text{as } S \to \infty, \tag{9}$$

where ϕ and Φ are the standard normal density and distribution function; as $S \to 0$, m^* tends to the stand-alone cost $\kappa[\rho(T) - \bar{p}]$ of Proposition 1.

(iv) Appetites. *If carrier j can hold at most $\bar{\theta}_j$ of the accumulation under its scenario appetite, cover also requires $\sum_j \bar{\theta}_j \geq 1$ at any price.*

(v) Autarky. *If instead each carrier keeps the accumulation of the policies it writes, a share w_j of the market, carrier j supplies cover if and only if $w_j/s_j \leq \theta_1^*$. With equal failure probabilities, the share of the accumulation written by carriers that exclude, $\sum_j w_j \mathbf{1}\{w_j/s_j > \theta_1^*\}$, does not fall when σ becomes more concentrated or when the accumulation grows relative to every s_j.*

When it can be shared, the accumulation is priced at the marginal capital of a book whose other risks are S times a full-market event, so that, to first order in $1/S$, the cost of carrying AI risk rises in proportion to H at any level of diversified capacity. This frictionless benchmark assumes that slices can be defined in contracts and traded and that carriers' other risks are independent of the provider factor. Exclusion arises where the accumulation stays with the carriers that write it, in books whose AI exposure is large relative to their other risks (part v) or whose appetites are small relative to the scenario (part iv), and the model takes this autarky as given. Not knowing which provider an insured relies on does not produce it on its own, because a reinsurer can charge for a slice what it would charge if all insureds used one provider, and that price would still be low except at the most intense exposures (Appendix C.3). More likely frictions are the price of reinsurance, since catastrophe reinsurance has been priced well above expected losses, which Froot (2001) attributes mainly to capital market imperfections and the market power of traditional reinsurers, and the definition of the slice, since liability wordings do not identify the provider behind a claim; for the same reason, a carrier that wants to shed the accumulation has to exclude AI as such. An event limit would define the slice contractually, but each claim would then have to be attributed to a provider event, although the insurer does not observe which model an insured used and harms may surface long after the defective update or reach an insured through a vendor's tool; Lloyd's required cyber wordings to settle the analogous attribution of attacks to states before such attacks could be excluded (Lloyd's 2022).

3.6 Where excluded losses go

In every state, the AI-related liability losses of adopters, $a\mu(\pi + qT)$, do not depend on c; coverage decides only who pays them. Insurers pay $c\,a\mu(\pi + qT)$ and firms retain the rest. Exclusion also creates deadweight default costs, proportional to $a\Psi(c)(\pi + qT)$ with $\Psi(c) = \Pr\left((1-c)X > B\right)$ decreasing in c, and credit losses for banks,

$$D(T) = \ell\, b\, a\, \Psi(c)\,(\pi + qT), \tag{10}$$

where b is lending per adopter and ℓ the loss given default. Credit losses inherit the structure of T: they jump in a year in which a provider fails, most of all when the largest one does.

Corollary 1 (Lenders' exposure)**.** *Let $p_k = p$ for all k, let carriers keep the accumulation of the policies they write, as in Proposition 2(v), and let σ' majorise σ.*

(i) The share of the accumulation that is excluded is no smaller at σ' than at σ.

(ii) For any coverage c, the systematic variance of banks' credit losses on loans to adopters and the rise in those losses when the largest provider fails,

$$\mathrm{Var}(D(T)) = \left(\ell ba\,\Psi(c)\,q\right)^2 p(1-p)H(\sigma) \quad \textit{and} \quad \ell ba\,\Psi(c)\,q \max_k \sigma_k,$$

are no smaller at σ' than at σ, and both rise when cover is withdrawn.

The internal ratings-based formula sees this exposure only through default probabilities: it applies the Basel formula, the 99.9% quantile of losses in a model with a single systematic factor (Gordy 2003; Basel Committee on Banking Supervision 2005), to banks' own estimates of default probabilities, which rise once AI-related defaults enter observed default rates. It therefore gives the same capital to a lender whose AI-using borrowers all rely on one provider as to a lender whose borrowers spread across providers. This is the point that Gordy (2003) and the literature on concentration risk (Pykhtin 2004; Düllmann and Masschelein 2007; Gordy and Lütkebohmert 2013) make about any second source of correlation, here a sector defined by the borrower's AI provider. Most of the formula's response is to idiosyncratic harms, whose defaults it treats as if they moved with the business cycle. When the provider factor is small and independent of the business cycle, the derivative of the 99.9% quantile of a model with both factors with respect to the added exposure is its expected value conditional on being at the quantile (Gouriéroux et al. 2000), close to its unconditional mean. The quantile therefore rises by about as much as expected loss does, and the unexpected loss of that model, the two-factor unexpected loss of Section 5, changes little. The provider factor shows up instead in the loss of a year in which a provider fails, which stress tests and assessments of concentration risk measure; Section 5 quantifies it and relaxes independence.

Upstream liability. Under the revised Product Liability Directive, a model developer is strictly liable to natural persons harmed by a defect, so part of the excluded loss can be recovered upstream. Recoveries are limited by what a provider can pay in the year its model fails. When that capacity does not grow with market share, the largest providers owe the largest claims when they can least pay them, and on the contracts as written the reform can raise the burden on the clients of large providers (Proposition 4 in Appendix G).

4 Quantification: the cost of carrying the accumulation

The model's most important parameters have never been measured: no liability-relevant failure of a foundation model at scale has been observed, and neither has the dependence of insureds and borrowers on particular models. I therefore treat the calibration as a set of stress scenarios and give the parameter values at which each conclusion reverses. Provider shares, adoption rates, insurance and banking aggregates, bank-level loan data and contract terms are data; the failure probability, the harm and default probabilities, insurers' appetites and, in Section 5.3, the share of a bank's corporate book lent to AI infrastructure and the loss rate on it are grids or illustrative values (Table 10), and Appendix C sets out the calibration and the results in full. Table 1 collects the main magnitudes. Risk measures are computed exactly by enumerating the 2^K failure states, and tables and figures use the stress case for the failure probability unless they state otherwise, so they lean towards exclusion and larger capital charges relative to the base case.

In brief, provider shares are the end-2025 shares, with a Herfindahl index of 0.28. The annual probability p that a provider has a failure relevant to liability is 0.4% in the base case and 1% in the stress case, from two output defects that reached interface customers in 13 to 15 provider-years. The premium-weighted adoption rate is 0.35, and the event intensity, the expected loss of a client of a failing provider in units of its annual liability premium, runs from 0.5 to 4 around a central value of 1, scaled on American claims practice; a harmed firm without cover defaults with probability 15% in the central case (Appendix C.1).

A liability book that mirrors the market carries the risk of the provider factor of at most 72 independent policies, however large it is. A failure of the largest provider costs 14% of a year's liability premium in the central scenario and 4 to 88% across the grid; with US liability cover silent on AI it would cost about USD 20 billion before reinsurance, 1.5% of the industry's surplus, a large loss for the liability lines but not a threat to the sector's solvency. Stand-alone capital under expected shortfall is close to the largest provider's share, 0.40 of a full-market event at $p = 1\%$ (Appendix C.2).

Assessed stand-alone, cover survives only if insureds' surplus from it, net of expenses, is at least 18% of expected losses in the base case and 22% in the stress case. Shared across the US property-casualty sector, the same accumulation carries an index of 0.01 to 0.02, less than a hundredth of the lower threshold, and its price rises roughly in proportion to the Herfindahl index; priced as though every insured relied on one provider, which a reinsurer that cannot observe the mix can charge, it carries 0.10 at the central exposure. Capital costs alone lead to exclusion only in books concentrated on AI-intensive insureds. Limits on the loss of a single scenario bind sooner: with an appetite of a tenth of equity, a liability writer with equity of 60% of premium breaches it once the AI share of its book times the intensity exceeds 0.15, while a carrier with the balance sheet of the US sector average breaches it only at an intensity of 6.5. The mechanism therefore bears on specialty and liability writers concentrated on AI-intensive insureds rather than on the multi-line groups that have filed most of the standard exclusions (Appendix C.3).

Concentration and exclusions have moved in opposite directions. With 2023 shares, when the largest provider held half of the market, the stand-alone index at $p = 1\%$ was 4.5 against 3.6 now, while the use of AI that generates language among large EU enterprises quadrupled; if accumulation plays a part in the exclusions, it does so through adoption. Limits per provider

Table 1. Headline magnitudes

Quantity	Base, $p = 0.4\%$	Stress, $p = 1\%$	Range	Basis
Herfindahl index of provider shares, end-2025 ($1/H$)	0.28 (3.6)		0.21–0.37; 0.49 in transaction data	Survey and transaction data
Loss from a failure of the largest provider, % of a year's liability premium	14		4–88; 20–21 with the other sources' largest shares	Adoption data; z assumed
Stand-alone exclusion index (thresholds 2.5 and 5)	3.0	3.6	1.9–6.6 with r; 5.4 with transaction shares	p and r assumed
Index in a book that mirrors the market, Euler (incremental)	0.16 (0.06)	0.37 (0.14)	up to 3.4 (2.1) on the grid	Other risks assumed
Index shared across the US property-casualty sector	0.01	0.02	0.04 with transaction shares; 0.10 with a single provider; 2.5 at a correlation of 0.34 with other risks	Premium data
Extra bank loss in the year the largest provider fails, euro-area significant institutions, % of CET1	0.33		0.2–1.3 over q and Ψ_0; 0.07–1.6 over z; peak year of payment 0.07	q and Ψ_0 assumed
Median EBA bank, % of CET1: borrowers mirror the market; all on the failing provider	0.28; 0.70		most exposed bank 0.9; 2.2; 22 banks above 1% in the second case	Loan data; q, Ψ_0 assumed
Share of corporate loans to clients of the failing provider at which a bank loses 1% of CET1, %: median ratio of loans to CET1; ninetieth percentile; highest	46; 27; 16		inversely proportional to $q\Psi_0$ (central 1.5%); about 35% of corporate loans go to AI adopters	Loan data; q, Ψ_0 assumed
Lender with CET1 of about a quarter of corporate loans, all adopters on the failing provider; specialist lender	0.8; 4.8		0.4–3.3; 2.4–19 over q and Ψ_0; 0.2–4.1; 1.0–23.5 over z	Illustrative

Notes: The base and stress cases differ only in the failure probability p. Other values: end-2025 shares, $q = 10\%$, $r = 0.1$, $a_w = 0.35$, $z = 1$, $\Psi_0 = 15\%$, other risks of a carrier with a standard deviation of 15% of its premium; ranges refer to $p = 1\%$ and run over the grids of Appendix C.1, and those of H over the survey snapshots and the splits of the residual share. The US property-casualty sector shares the accumulation across USD 1.1 trillion of premium, with other risks independent of it unless a correlation is stated. Bank figures assume that none of the AI-related liability of borrowers is insured and scale with the share of borrowers whose liability is excluded; ranges over z tie Ψ_0 to the intensity (Table 18), and the peak year of payment uses the payout pattern of US other liability. Shares of corporate loans are evaluated at $q\Psi_0 = 1.5\%$ for the banks with the median, the ninetieth-percentile and the highest ratio of loans to non-financial corporations to CET1 capital. The concentrated lender holds the ratio of CET1 capital to corporate loans of euro-area significant institutions, about a quarter, and 35% of its corporate loans go to AI adopters; the specialist lender holds CET1 capital of an eighth of its loans, all to clients of the failing provider.

event relax both constraints: a limit of a tenth of a full-market event per failing provider lowers the stand-alone index from 3.6 to 1.0 while 95% of the expected AI loss remains insured (Appendix F).

5 Where excluded losses land: banks

This section computes what the exclusion of AI liability would mean for lenders, first for stylised portfolios and then bank by bank. All figures assume that none of the AI-related liability of borrowers is insured and scale in proportion to the share of borrowers whose AI-related liability is excluded. That proportionality holds for the liability channel of Sections 5.1 and 5.2; the asset channel added in Section 5.3 runs through lending to the AI industry itself

Table 2. Losses on corporate loans to AI adopters under full exclusion

			Failure-year	Change under exclusion		
Portfolio	q	Ψ_0	extra loss	EL	Two-factor UL	IRB UL
Idiosyncratic and common harms ($\pi = 0.1q$)						
Corporate	5%	15%	12	3.3	−0.4	17.9
Corporate	10%	15%	24	6.5	−0.8	34.5
Corporate	10%	30%	48	13.1	−1.6	64.6
Corporate	20%	30%	96	26.1	−2.8	115.5
Provider failures alone ($\pi = 0$)						
Corporate	5%	15%	12	0.3	0.0	1.7
Corporate	10%	15%	24	0.6	0.0	3.3
Corporate	10%	30%	48	1.2	0.0	6.6
Corporate	20%	30%	96	2.4	0.4	13.1

Notes: Basis points of exposure at default. Failure-year extra loss = LGD $\times\, q\,\sigma_{\max}\Psi_0$ when the largest provider fails. The last three columns give the change under exclusion in expected loss, in the two-factor unexpected loss (the 99.9% quantile less expected loss in a model with the macroeconomic factor and provider failures, independent of each other) and in unexpected-loss capital under the internal ratings-based formula, without maturity adjustment, when the default probability of 1% is re-estimated to include AI-related defaults. $p = 1\%$ (the stress case), end-2025 shares, LGD = 40%; SME portfolios are in Table 19 in Appendix C.

and does not scale with the exclusion share, so the combined figures reported there do not either.

5.1 Portfolio losses and Pillar 1 capital

With idiosyncratic and common harms, full exclusion raises expected credit losses on loans to AI adopters by 3 to 26 basis points a year, nine-tenths of them from idiosyncratic harms that insurers used to pay, and in a year in which the largest provider fails, losses rise by a further $\ell q \sigma_{\max}\Psi_0$, 12 to 96 basis points (Table 2). Once banks re-estimate default probabilities to include AI-related defaults, unexpected-loss capital under the internal ratings-based formula rises by 18 to 115 basis points for corporate portfolios and by 7 to 49 for SME portfolios, mostly for the idiosyncratic harms, because the formula gives their defaults the same dependence on the business cycle as any other default (Gordy 2003), while the two-factor unexpected loss, computed with both the business cycle and provider failures, falls slightly. For provider failures alone, the add-on is 2 to 13 basis points for corporate and 1 to 5 for SME portfolios, and the two-factor unexpected loss changes by less than half a basis point. Pillar 1 capital therefore spreads excluded AI losses across lenders without regard to concentration, and it does so only once default estimates embed AI-related defaults, which takes years of default history. With Ψ_0 tied to the intensity, the failure-year loss of a lender whose borrowers mirror the market ranges from 5 basis points at $z = 0.5$ to 118 at $z = 4$ (Table 18 in Appendix C).

Lenders differ in the concentration of their borrowers' dependence, which the formula does not see. A lender whose AI-adopting borrowers all rely on the failing provider loses 30 to 240 basis points of that exposure in the failure year, $1/\sigma_{\max} = 2.5$ times as much as a lender whose borrowers mirror the market, while the two-factor unexpected loss of its corporate portfolio changes by between −1 and 1 basis point, because a failure that harms at most a fifth of its clients moves default rates far less than a severe downturn. With CET1 capital of about

a quarter of its corporate loans, as for euro-area significant institutions in aggregate, and 35% of those loans to adopters, such a lender would lose 0.4 to 3.3% of its CET1 capital. A lender specialised in technology firms, with all its borrowers on the failing provider and a balance sheet of corporate loans at full risk weight, so that a CET1 ratio of 12.5% means CET1 capital of an eighth of its loans, would lose 2.4 to 19% of its CET1 capital, and 23.5% at $z = 4$ (Silicon Valley Bank, whose clients were concentrated in venture-backed technology and life-science firms, reported a CET1 ratio of 12% at the end of 2022; Board of Governors of the Federal Reserve System 2023). With about 0.6 percentage points of annual loan growth associated with each percentage point of the ratio of equity to assets (Gambacorta and Shin 2018), such a lender would expand its lending by 0.1 to 1.0 percentage points less in the year after the failure if its loans are two-thirds of its assets. Because liability claims are paid over several years, the peak annual loss would be about a fifth of these failure-year figures if defaults follow payments, which in the Schedule P triangles for US other liability peak at 21% of an accident year's ultimate loss in the third year (Meyers and Shi 2011).

5.2 Bank by bank

Of the 108 banks of the EBA's 2025 transparency exercise that report their loans to non-financial corporations by NACE section for June 2025, I exclude four whose risk-weighted assets are less than a tenth of their total assets, against 36% at the median bank. Risk weights that low indicate loans to or guaranteed by the public sector, to which a loss given default of 40% does not apply. The remaining 104 banks, from 25 EU and EEA countries, together with one bank whose country is not disclosed, held EUR 6.33 trillion of such loans and EUR 1,548 billion of CET1 capital (European Banking Authority 2025). I scale the central adoption rate of 35% by each bank's sector-weighted share of enterprises that use AI generating language (Eurostat 2026b), relative to the same share for all banks. Sector composition has little effect: loans to information and communication and to professional services, the sections that use these tools most, are 8% of corporate loans at the median bank and 15% at the ninetieth percentile, and the relative adoption rate lies between 0.83 and 1.11 for four-fifths of the banks. A bank's loss therefore depends mainly on the ratio of its corporate loans to its CET1 capital, from 1.43 at the tenth percentile to 6.26 at the ninetieth, which accounts for nearly all the variation in computed losses across banks, and on the share of those loans that go to clients of the failing provider, which is unobserved (Table 20 in Appendix C). At $q = 10\%$ and $\Psi_0 = 15\%$, the median loss across banks would be 0.28% of CET1 capital if borrowers mirror the market, 0.37% if borrowers in information and communication and professional services all rely on the failing provider, and 0.70% if all adopting borrowers do. With borrowers that mirror the market no bank would lose 1%, and the most exposed would lose 0.87%; with all adopting borrowers on the failing provider, 22 of the 104 would lose more than 1%, and the most exposed bank 2.2% (Figure 11b in Appendix H). Against the headroom above the average CET1 requirements and guidance of significant institutions for 2026, 11.2% of risk-weighted assets (European Central Bank 2025), the median share of headroom used across banks would be 0.8% in the first case and 2.1% in the last; no bank would breach it, and two would use more than a tenth of it.

Run backwards, the bank with the median ratio of corporate loans to capital loses 1% of its CET1 capital only if 46% of its corporate loans go to clients of the failing provider, more

than the share that goes to AI adopters at all, while the bank with the highest ratio gets there at 16% (Appendix H.2). Euro-area significant institutions together would lose EUR 5 billion, or 0.33% of their CET1 capital, and no bank would exhaust its buffer above requirements (Appendix H.3). For the system as a whole, the case for supervisory attention rests on lenders whose borrowers' dependence is concentrated.

5.3 Both sides of the balance sheet

Banks hold the provider factor a second time, as lenders to the AI industry itself, and this is the channel to which supervisors have begun to turn. In the ECB's bank lending survey of October 2026, several banks expected a gradual but sustained increase in demand for loans to finance digital infrastructure, particularly data centres, and banks reported that their current and prospective AI-related lending is more likely to target infrastructure providers than software developers or firms adopting AI (European Central Bank 2026c). The ECB has reportedly asked a small number of banks for details of their lending to data centres (Duarte et al. 2026), and its Financial Stability Review notes that private credit investors in AI and software firms, unlike private equity and venture capital investors, bear the losses from failures without sharing in the successes (European Central Bank 2026a). Private credit financed 34% of AI investment in 2025, against 9% in 2024, on OECD estimates cited by the Bank of England (Bank of England 2026), so part of the exposure reaches banks as lenders to those funds rather than to the projects. Part of this debt is serviced from long-term contracts with the same few providers: in July 2025 a cloud provider closed a USD 2.6 billion term loan from ten banks, three of them from the euro area, secured by the assets of a subsidiary set up to serve a long-term agreement with OpenAI (DatacenterDynamics 2025).

The two exposures are not independent. A failure severe enough to impair a dominant provider's revenue reaches the infrastructure lent against it in the same state in which its clients' excluded liabilities reach the deployers' lenders, so a bank exposed to both takes both losses in that state and equation (11) adds them. That addition is an accounting identity within the scenario and not a statement about the joint distribution. The failure modes differ, and that limits how far the scenario travels: a defect that produces claims need not impair revenue, and a regulatory suspension impairs revenue without producing claims (Table 24), so the combined case is the severe one in which both follow rather than the typical incident, and the probability of that conjunction is not estimated here. Only where an asset loss is non-decreasing in T are the two comonotonic, so that expected shortfall adds them without diversification (Dhaene et al. 2006); that is the case treated for insurers in Appendix F.

Write ι for the share of a bank's corporate loans lent to AI infrastructure and ϱ for the loss rate on that exposure in the failure year. The combined failure-year loss is then

$$\frac{\text{loss}}{\text{CET1}} = \frac{L}{K}\Big(\underbrace{x\,\text{LGD}\,q\Psi_0}_{\text{liability}} + \underbrace{\iota\,\varrho}_{\text{asset}} \Big), \tag{11}$$

with L the bank's loans to non-financial corporations and K its CET1 capital. I take $\varrho = 12.5\%$ centrally, in a range of 5 to 25%: the product of a default rate on facilities whose contracted revenue comes from the failing provider and a loss given default on lending secured by data centres and processors, whose collateral value depends on the same demand and depreciates

Table 3. Combined failure-year loss, liability and asset channels

		Liability	Combined, by ι		
Bank, by NFC loans / CET1	L/K	only	1%	2%	5%
Borrowers mirror the market					
Tenth percentile	1.43	0.12	0.30	0.48	1.01
Median	3.63	0.30	0.76	1.21	2.57
Ninetieth percentile	6.26	0.53	1.31	2.09	4.44
Highest	10.48	0.88	2.19	3.50	7.43
Adopting borrowers all on the failing provider					
Tenth percentile	1.43	0.30	0.48	0.66	1.19
Median	3.63	0.76	1.22	1.67	3.03
Ninetieth percentile	6.26	1.32	2.10	2.88	5.23
Highest	10.48	2.20	3.51	4.82	8.75

Notes: Per cent of CET1 capital in the year the largest provider fails, from equation (11), for the 104 banks of Table 20 ranked by the ratio L/K of loans to non-financial corporations to CET1 capital. ι is the share of the corporate book lent to AI infrastructure and the asset channel is $(L/K)\,\iota\,\varrho$ with $\varrho = 12.5\%$. Liability channel at LGD $= 40\%$, $q\Psi_0 = 1.5\%$ and an adoption rate of 35%, with $x = 0.35 \times 0.40$ where borrowers mirror the market and $x = 0.35$ where adopting borrowers all rely on the failing provider. Inverting the combined loss in the concentrated case, the bank with the median ratio reaches 1% of CET1 capital at $\iota = 0.52\%$ (0.26 at $\varrho = 25\%$, 1.31 at $\varrho = 5\%$), and the banks at the ninetieth percentile and above reach it through the liability channel alone.

quickly. Like q and Ψ_0, ϱ is a parameter of the scenario and not an estimate.

At modest exposures the asset channel is the larger of the two (Table 3). It matches the whole liability channel once AI infrastructure reaches 1.7% of a bank's corporate book where all of its adopting borrowers rely on the failing provider, and 0.7% where its borrowers mirror the market. The bank with the median ratio of corporate loans to capital, which the liability channel alone never takes to 1% of CET1 capital, passes that mark with 0.5% of its corporate book in AI infrastructure, and at 2% its combined loss is 1.7%. For the bank with the highest ratio the combined loss is 4.8% at $\iota = 2\%$ and 8.8% at 5%, against 2.2% from excluded liability alone.

The firms that build this infrastructure file with the SEC. Ten US-listed AI-infrastructure firms disclose 31 financings since 2025, USD 49.9 billion in all (Table 13 in Appendix C). JPMorgan Chase and Morgan Stanley arrange eleven of the sixteen loan facilities that have a bank as agent, and they are also agents to 15 of the 31 facilities of firms that name a provider, so the credit on both sides of equation (11) is arranged by the same houses. The credit agreements that some of these borrowers file tie their lending to named providers (Section 6.5), but no filing says how much of a facility each bank retains and no supervisor publishes this lending by bank, so ι remains a parameter of the scenario (Appendix H.4).

6 What can be measured: dependence, lenders and prices

Sections 5 and 5.3 rest on two quantities that no supervisor publishes, the share of a lender's book that goes to clients of one provider and the share lent to infrastructure serviced by the same providers, and Section 2.1 traced the chain from deployers to their lenders only as far

as the administrative agent of each facility. This section measures what public filings and commercial loan data allow: which firms say that their products run on which provider's models, which lenders hold their syndicated loans and in what shares, whether the shares of these firms respond when their provider has an incident, whether lenders, bondholders and the firms themselves act on that response, and what the contracts with providers say about who bears a failure. The code and the hand coding are in the replication files.

6.1 Dependence in annual reports

A full-text search of annual reports filed with the SEC on Forms 10-K, 10-KT, 20-F and 40-F, and their amendments, from January 2023 to early November 2026 (U.S. Securities and Exchange Commission 2026), for phrases that name a model provider (OpenAI, ChatGPT, Azure OpenAI, Anthropic, Google Gemini, Vertex AI, Meta Llama, Mistral AI, Cohere and Amazon Bedrock) returns 571 documents from 326 firms. From each document I keep the passages around every mention, 450 characters on either side, which leaves 541 documents from 313 firms with at least one passage. I then code each firm by hand from its passages. Eighty-nine firms state that a product or service they sell runs on a provider's model; 25 use one only internally, for example through enterprise licences for staff; 13 supply or invest in providers, among them chip makers, cloud operators and content licensors; three are providers themselves. The remaining 183 mention a provider only as an example of the technology, as a competitor, in litigation, in investment-fund boilerplate or in a director's biography. Thirty-five of the 89 are foreign private issuers, and 47 first named their provider in a report filed in 2023 or 2024.

Disclosed dependence is far more concentrated than spending (Table 4). Seventy-three of the 89 firms name OpenAI, 56 of them no other provider; 20 name Google, 15 Anthropic and 21 more than one. Splitting each firm equally across the providers it names gives OpenAI 71% of disclosed product dependence and a Herfindahl index of 0.54, against 0.28 for end-2025 spending and the range of 0.21 to 0.37 across the survey snapshots of Section 2.1. The difference has three sources that the data cannot separate. Disclosure lags adoption, and OpenAI took half of enterprise spending in 2023, when many of these products were built; OpenAI's consumer product is the name most firms reach for; and firms that build on a provider's models through Microsoft's cloud are counted under OpenAI. The disclosures therefore identify too few dependent firms, since most adopters name no provider (Section 2.1), and possibly too concentrated a set, and I do not use them to replace the spending shares of the calibration. What they add is a list of firms whose dependence is on record, which can be followed to their lenders and to their share prices.

6.2 Who holds the loans

DealScan records syndicated loans with every lender in the syndicate and, for part of them, the share each holds (LSEG 2026). I match the coded firms to DealScan borrowers by ticker and then by name, and check every match by hand, which discards seven false matches on tickers shared across exchanges. Thirteen of the 89 firms whose products run on a provider's model and 12 of the 25 that use one internally borrow in the syndicated market; most of the rest are small software firms without syndicated debt. To the infrastructure side I add the firms

Table 4. Providers named by firms that disclose dependence in annual reports, 2023–2026

	Products built on the model			Products or internal use	
Provider named	Firms	Only provider	Share (%)	Firms	Share (%)
OpenAI	73	56	71.4	93	72.9
Anthropic	15	3	9.3	18	9.0
Google	20	8	14.9	23	13.8
Meta	3	0	1.5	3	1.2
Mistral	1	0	0.3	2	1.1
Cohere	2	1	1.5	2	1.2
Amazon (Bedrock)	2	0	1.1	2	0.9
Firms	89			114	
Herfindahl index			0.54		0.56

Notes: Firms whose annual reports on Forms 10-K, 10-KT, 20-F or 40-F filed from January 2023 to November 2026 name a provider and state that a product or service they sell runs on its models (first three columns), or that they use its models in products or internally (last two). "Only provider": firms that name no other. Shares split each firm equally across the providers it names. Azure OpenAI is counted under OpenAI, and Amazon Bedrock is a hosting service through which several providers' models are sold.

of Section 5.3 that DealScan records: CoreWeave and two of its financing vehicles, Galaxy Digital's Helios data-centre subsidiary and TeraWulf. DealScan reports a lender's share for 18% of lender-tranche records; where it is missing, the part of a tranche that reported shares leave unallocated is divided equally among the other lenders, and amounts are commitments, drawn or not. Exposure is measured at 30 June 2026, over tranches active then and not yet matured.

At that date 22 borrowers that disclose dependence had USD 50.0 billion of syndicated facilities outstanding, held by 86 lenders, and the four infrastructure borrowers USD 27.6 billion, held by 30 (Table 27 in Appendix H). Every one of the 25 largest syndicated lenders holds some of the first, with a median of 0.41% of its syndicated book and a maximum of 0.93%, and at most 1.5% once infrastructure is added; one private-credit lender holds 2.7%. These shares are far below the 16 to 46% of corporate loans that the reverse stress test of Appendix H.2 requires for a loss of 1% of CET1 capital, but they are floors, because only firms that name their provider are counted, and the denominators differ, a syndicated book against all corporate loans. What the data show more firmly is where the exposure sits. The five largest holders have 27% of the lending to dependent firms, against 17% of all syndicated lending, and the Herfindahl index across lenders is 0.033, against 0.014; on the infrastructure side the five largest hold 31%. The two channels are held by the same institutions: 19 of the 30 infrastructure lenders also lend to dependent firms and hold 59% of that lending, and six of the ten largest infrastructure lenders are among the ten largest lenders to dependent firms. Equation (11) adds the two channels within a scenario; these holdings show that the addition happens on the same balance sheets.

6.3 What share prices register

If dependence on one provider is a common factor, the clients of a provider should move together when that provider has an incident, and apart from the clients of its competitors. I test this on the dated incidents of Table 24 at OpenAI (with incidents at Microsoft's Azure OpenAI

service counted under OpenAI), Google, Anthropic, Meta and Mistral. For each incident I compare the abnormal returns of firms whose products run on the affected provider's models with those of firms whose products run only on other providers' models, counting a firm from twelve months before its first annual report that names the provider. Abnormal returns come from a market model estimated over trading days -250 to -30 against the S&P 500, from adjusted closing prices of 80 of the 89 firms (Yahoo Finance 2026). I keep the 38 incidents with at least three firms on each side, 25 at OpenAI, seven at Anthropic and six at Google, which give 1,455 observations of a firm exposed to an incident. Inference uses one difference per incident, which absorbs the correlation among firms hit on the same day.

Over the incident day and the next, the clients of the affected provider fall by 1.42 percentage points relative to the clients of its competitors, with a t of -3.6 (Table 5 and Figure 3). The difference is -0.81 points on the incident day itself, -1.47 from the day before to the day after, and -0.95 over four days. None of 1,000 sets of random dates, drawn away from any incident and assigned to the same providers, gives a difference as negative; the fifth percentile of the placebo distribution is -0.62 points. The result holds against a software-sector benchmark, with each firm's two-day abnormal return capped at 15%, when any single incident is dropped (between -1.29 and -1.57 points), for outages and for the other types of incident, and for 2025 and 2026 alone. It is a reallocation more than a fall: the clients of the affected provider lose 0.90% on average ($t = -2.2$) while the clients of its competitors gain 0.52%. At OpenAI's incidents the difference comes from rivals' clients rising, at the other providers' incidents from the affected clients falling, by 2.7%.

Naming more than one provider does not protect a firm. Nineteen of the 80 firms name several providers. Relative to the clients of other providers, firms that name only the affected provider fall by 1.37 percentage points over days 0 and 1 ($t = -2.8$), and firms that name it among several fall by 1.67 points ($t = -3.8$); the gap between the two, 0.26 points in the other direction, is far from significant ($t = 0.5$ across 37 incidents). Firms that contract with several providers may be the most AI-intensive, or may run their main products on one of them. Section 6.4 follows firms as they add providers and asks whether doing so changes their reaction.

Two qualifications bound what this shows. The reaction does not depend on whether a firm had yet disclosed its provider: the affected provider's clients fall by about as much before their first annual report that names it as after, so investors appear to learn of dependence from product announcements and other sources before annual reports record it, and the test identifies dependence rather than an effect of disclosure. And the firms are mostly small, listed software companies, whose shares react to news that a large, diversified borrower's would absorb. Within these limits, equity prices register the provider factor that the instruments of Sections 5 and 2.5 do not: a common movement among the clients of one provider, of about one and a half percentage points over two days for incidents most of which lasted hours or days.

6.4 Lenders' prices and firms' sourcing

Two parties could act on the factor that share prices register. Lenders could charge for it, and the firms themselves could reduce it by adding providers. The data allow a test of each.

DealScan records the all-in spread over the base rate for about a third of tranches. I regress

Table 5. Abnormal returns of clients around provider incidents, by provider

	Clients of the provider		Others	Difference		Incidents
	CAR	(t)	CAR	CAR	(t)	
Panel A. Windows, market model against SPY						
Day 0	−0.23	(−0.8)	0.59	−0.81	(−2.5)	38
Days 0 to 1	−0.90	(−2.2)	0.52	−1.42	(−3.6)	38
Days −1 to 1	−0.77	(−1.5)	0.71	−1.47	(−3.7)	38
Days 0 to 3	−0.67	(−1.2)	0.27	−0.95	(−2.0)	38
Panel B. Days 0 and 1, variations						
IGV benchmark	−0.34	(−0.9)	1.09	−1.42	(−3.6)	38
CARs capped at 15%	−1.07	(−3.2)	0.11	−1.18	(−4.0)	38
OpenAI and Microsoft incidents	0.03	(0.1)	1.20	−1.17	(−2.1)	25
Other providers' incidents	−2.68	(−3.5)	−0.80	−1.88	(−4.2)	13
Outages	−0.87	(−1.1)	0.83	−1.70	(−2.7)	16
Other incident types	−0.92	(−2.1)	0.29	−1.21	(−2.4)	22
2025–2026 only	−1.03	(−2.0)	0.27	−1.30	(−3.8)	26
Clients naming only that provider	−0.89	(−1.7)	0.47	−1.37	(−2.8)	37
Clients naming several providers	−1.15	(−2.6)	0.52	−1.67	(−3.8)	38
Provider sets dated by filing	−1.02	(−2.3)	0.54	−1.56	(−3.8)	37
Same provider, before disclosure	−0.90	(−1.3)	−1.20	0.29	(0.7)	15
Placebo: 1,000 sets of random dates; share with a difference at least as negative: none						

Notes: Cumulative abnormal returns (CAR) in per cent, from a market model estimated over trading days $[-250, -30]$, averaged over firms within each incident and then across incidents; t-statistics across incidents in parentheses. Clients of the provider: firms whose annual reports state that their products run on the affected provider's models; others: firms whose products run only on other providers' models. In the rows on clients naming only that provider or several providers, the clients of the affected provider are restricted accordingly. In the row "Provider sets dated by filing" each provider counts from twelve months before the first annual report that names it, including providers added in later reports (Table 28 in Appendix H). In the row "Same provider, before disclosure" the comparison group is instead the affected provider's clients before their first annual report that names it, less twelve months. Incidents with at least three firms in each group. Placebo: the mean difference over days 0 and 1 for the same providers on random dates at least six days from any incident, with abnormal returns from a market model over the whole sample.

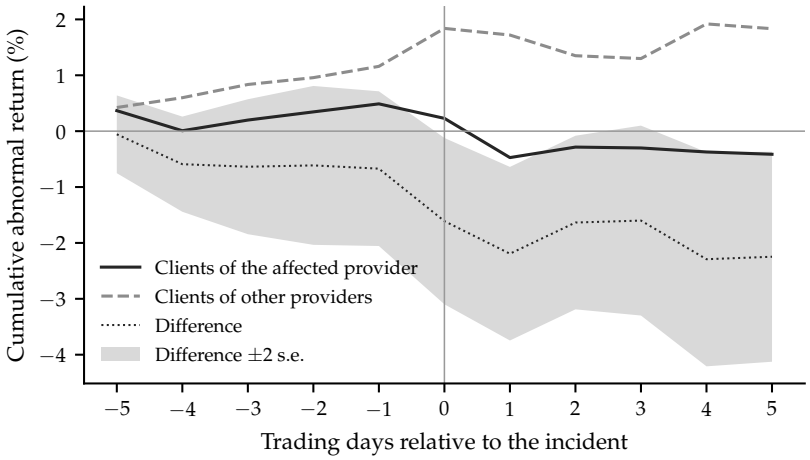


Figure 3. Clients of the affected provider and of its competitors around provider incidents
Notes: Mean cumulative abnormal returns across the 38 incidents of Table 5, cumulated from five trading days before the incident, and their difference with a band of two standard errors across incidents.

the spread of 25,790 tranches to 11,191 borrowers made from 2021 to 2026 on indicators for the tranches of borrowers that disclose dependence, split by whether the tranche was made after the date from which the borrower counts as dependent in Section 6.3 or before it, and for the infrastructure borrowers of Section 6.2. The regressions control for loan terms and for industry-by-year and country effects, and standard errors are clustered by borrower (Table 6). Seventeen borrowers that disclose dependence have priced tranches, 27 made after the date and 23 before. Lenders charge them no premium. Tranches made after the borrower counts as dependent are priced 7.5% below comparable loans, with a standard error of 12.3%, so the confidence interval excludes a premium above about a sixth of the spread, or about 43 basis points; tranches made before the date are priced the same way, and the difference between the two is nil ($t = 0.0$). The result is the same for borrowers in North America and Europe, and within the seven borrowers that took loans on both sides of the date, whose spread rises by an insignificant 9.8% (standard error 20.3%). Infrastructure borrowers, by contrast, pay 35 log points, or 87 basis points, more than comparable loans ($t = 2.9$). That premium prices young, leveraged borrowers whose loans are secured on chips and on contracts with providers, and the regression cannot separate it from their credit quality; what it does show is that the absence of a premium for dependent borrowers is not a failure of the data to register pricing at all. The test has limits of its own: spreads are observed only when a loan is made, the sample can rule out large premiums but not small ones, and the borrowers are larger and more diversified than the firms whose shares react most. Within these limits, lenders price the loans of firms whose shares move with their provider as if the provider did not matter, which is what credit files that do not record the provider would lead them to do.

Bond markets give a second test, on firms of the size whose credit banks hold. TRACE records institutional trades in the bonds of 51 of the firms that name a provider, 16 of them firms that disclose dependence (FINRA 2026); the data available to me end in March 2026, which leaves 31 incidents. Around each incident I compare the change in bond prices from the last trade in the three days before day 0 to the first trade in the three days after it, averaged over each issuer's bonds, between the dependent clients of the affected provider and the other issuers. The difference is -0.04 percentage points for prices (standard error 0.10) and $+0.6$ basis points for yields (standard error 4.6), with about five treated issuers and 29 others per incident. The bonds of firms whose shares fall by about one and a half points do not move, although the test has little power: a fall of up to a quarter of a percentage point cannot be excluded. Yield levels point the same way. With bond and month effects, the yields of dependent issuers change by -17.9 basis points (standard error 14.0) after the date from which they count as dependent, identified from the 14 issuers with bonds trading on both sides of it, so a rise of more than about ten basis points is excluded. As far as these data reach, neither lenders nor bondholders price the factor that equity investors register.

Firms can also reduce their dependence by adding providers, and their annual reports show whether they do. Following the firms that disclose dependence through their successive reports, and counting a provider when a report names it outside a description of competitors or of the market (Table 28 in Appendix H), the share of reports that name more than one provider rises from none in 2023 to 8% in 2024, 15% in 2025 and 31% in 2026. The change comes from new reports more than from firms changing their sourcing. In 88 pairs of consecutive reports by firms that already named a provider, I check every change by hand: ten firms add a

Table 6. Loan spreads of borrowers that disclose dependence on a provider

	(1) Log spread All	(2) Spread (bp) All	(3) Log spread N. America, Europe	(4) Log spread Repeat borrowers
Dependent, after dependence	−0.075	−7.5	−0.070	0.098
	(0.123)	(25.7)	(0.128)	(0.203)
Dependent, before dependence	−0.068	−23.2	−0.065	
	(0.100)	(22.9)	(0.100)	
Infrastructure borrower	0.350	86.6	0.327	
	(0.120)	(34.8)	(0.120)	
Loan terms	yes	yes	yes	yes
Industry-year, country effects	yes	yes	yes	year only
Borrower effects	no	no	no	yes
Tranches	25,790	25,790	18,704	20,735

Notes: Ordinary least squares on syndicated tranches made from 2021 to 2026 with an all-in drawn spread in DealScan (LSEG 2026); standard errors clustered by borrower in parentheses. Dependent borrowers: borrowers that state in annual reports that their products run on a provider's models or that they use one internally (Section 6.2); a tranche counts as after dependence if made on or after the date from which the borrower counts as dependent in Section 6.3, twelve months before its first annual report that names the provider. Infrastructure borrowers: CoreWeave and its financing vehicles, Galaxy Digital's Helios subsidiary, Hut 8, Nebius and TeraWulf. Loan terms: log tranche amount, maturity, investment-grade and leveraged market segments, security, sponsor backing, log number of lenders and tranche type. Column (4) absorbs borrower effects on the borrowers with at least two priced tranches, so the dependence indicator is identified from the seven dependent borrowers with tranches on both sides of the date.

provider and two drop one. None of the additions appears in the reports filed in 2024, the first after OpenAI's governance crisis of November 2023 put its continuity in doubt; they come in 2025 and 2026, as enterprise spending became less concentrated on OpenAI (Section 2.1). And every firm that adds a provider keeps the one it had. Firms diversify by layering a second provider on the first, slowly, not by switching away after a provider's failure.

Layering does not appear to change how the market treats a firm. Dating each firm's providers by the reports that name them, rather than by the firm's full list, leaves the main result unchanged, at −1.56 percentage points ($t = -3.8$; Table 5). With these dated sets, 24 of the 80 firms name several providers at the time of some incident, and the clients that name only the affected provider and those that name it among others fall by the same amount (a gap of 0.02 points, $t = 0.0$ across 35 incidents). Within firms, with firm effects, the reaction is 1.46 points smaller in incidents after a firm names a second provider, but the estimate rests on the eight firms whose status changes and is not significant ($t = 0.6$). The data therefore cannot show that adding a provider takes a firm out of its first provider's factor. They show that firms add providers slowly, keep the first one, and move with it at the dates of its incidents whether or not they have added others. This supports the assumption of Section 3.1 that each adopter relies on one provider: a firm that adds others still moves with its first, and its own sourcing does little to remove the factor of Lemma 1.

6.5 Contract terms

The evidence so far is about prices. Contracts show who is meant to bear a provider's failure, and two kinds of filing reveal them: what dependent firms tell investors about their terms

with providers, and the agreements with providers that suppliers of computing capacity file as exhibits. I read the full annual reports of the 114 firms that disclose dependence, 228 reports in all, keep every paragraph on AI and providers, and code what they say about their terms (Table 7, Panel A). I then search the full text of EDGAR filings since 2021 for the legal names of the providers and read every agreement filed as an exhibit in which one of them is a party (Panel B).

Deployers describe the dependence but not the protection. Sixty-five of the 114 firms name a provider or third-party models in their risk factors, 27 of them as a dependence, and they name the ways it can fail: outages (23 firms), discontinuation or loss of access (19) and changes in price or terms (26). Ten describe an architecture that uses several models, and four say that switching provider would be slow, costly or uncertain. Contractual recourse barely appears. Three firms mention indemnities from providers, each to say that they may not cover its losses, and one mentions providers' limits of liability; none reports a service-level commitment from a provider, and none reports insurance that covers AI output or a provider's failure. This matches the published terms of Section 2.3, which leave deployers little recourse. In annual and quarterly reports, registration statements and current reports filed from 2021 to 2026, the phrases "AI exclusion", "affirmative AI", "silent AI" and "AI liability insurance" do not occur; of the 36 documents that mention "AI insurance", the nine whose passages could be extracted all use it for AI applied to insurance. The exclusions of Section 2.4 are thus absent from what insured firms disclose, which is where a lender's credit analysis would look for them.

The agreements filed by suppliers of computing capacity show where protection goes instead. Under CoreWeave's master services agreement with OpenAI, OpenAI pays fees for reserved capacity and the published text allows termination only for uncured breach or insolvency; if CoreWeave becomes insolvent, OpenAI may require the contract to be assigned to a substitute operator of its choice, and OpenAI gives CoreWeave no credit support (CoreWeave, Inc. 2025d). Cerebras's agreement commits OpenAI to 250 MW of capacity by the end of 2026 and 500 MW by the end of 2027 and protects OpenAI against late delivery by shortening the term (Cerebras Systems Inc. 2026). Anthropic's order with Nscale pays fees net of service credits, requires the supplier to indemnify and insure, and obliges it to finance the chips and the data centre on terms set for a customer with Anthropic's credit profile, without Anthropic taking on any liability in that financing (Nscale Ltd 2026); Akamai's agreement makes service credits Anthropic's only remedy for interruptions (Akamai Technologies, Inc. 2026). The only support a supplier receives comes from a third party and is optional: Google may, but need not, cure Anthropic's defaults under its lease from a TeraWulf subsidiary, or take the lease over (TeraWulf Inc. 2026).

The suppliers' lenders take the provider's credit directly. The facility that finances CoreWeave's capacity for OpenAI at one site lends to a special-purpose borrower whose revenue is the OpenAI contract; cash is trapped if OpenAI breaches or fails to pay for three months, and CoreWeave's own guarantee falls away once OpenAI is rated investment grade by S&P and Moody's (CoreWeave, Inc. 2025b). Another facility requires a coverage ratio of at least 1.40 on projected contracted cash flows from Microsoft and OpenAI (CoreWeave, Inc. 2025c), and CoreWeave's revolving facility permits debt secured on the OpenAI contract up to 75% of its value, or 90% once OpenAI is rated investment grade, and up to USD 7 billion (CoreWeave, Inc. 2025a). Banks also lend to the providers themselves: the secured revolving

facility that Anthropic amended in September 2026 has 17 arranging banks, five of them European (Anthropic, PBC 2026).

Taken together, the contracts place the provider factor where Sections 5 and 2.5 put it. Providers protect themselves against their suppliers' failure; deployers report the dependence but neither recourse nor cover; and the suppliers' lenders hold the provider's credit under covenants that turn on its payments and its rating, in loans that credit files record as exposures to CoreWeave or to a special-purpose vehicle. The covenants show that lenders know whose credit they are taking. What no published data show is how much of it each bank holds, which is the quantity Section 7 asks supervisors to collect.

7 Policy

Recording dependence. Neither underwriting nor credit files say which models an insured's or a borrower's operations depend on, and annual reports rarely name a provider (Section 2.1). The bank-level data show what the gap hides: sector composition barely moves a bank's exposure, while the concentration of its borrowers' dependence can multiply its failure-year loss by up to $1/\sigma_{\max}$, two and a half at end-2025 shares. In cyber, the requirement to identify non-affirmative exposure led the market to build exposure data and disaster scenarios that had not existed before (Prudential Regulation Authority 2017; Lloyd's 2019). The analogue for AI is a field, collected at underwriting and in credit files for large exposures, that records whether core processes depend on a foundation model, on which one, and whether a workflow has a fallback. The same field is needed on the other side of the balance sheet, where it would record which provider's contracts service the data centres and other AI infrastructure a bank has lent against; the filings of Section 5.3 show that some facilities name that provider and others do not, and nothing requires them to. Aggregated by supervisors, such fields would measure the concentration of dependence and identify the lenders and carriers in which it is concentrated, and disclosure guidance for annual reports could raise the share of firms that name their provider. Section 6 shows how far existing sources reach: commercial loan data record syndicates and most shares, and annual reports name the provider for a minority of firms, so the missing piece is the link between the two for the many firms that name none. A cheaper first step is at hand: the ECB's bank lending survey has already put ad hoc questions on AI-related lending to banks (European Central Bank 2026c), and an ad hoc question on the shares of corporate exposures whose borrowers' core processes rely on each of the main providers, and on exposures to AI infrastructure whose contracted revenue comes from each of them, would let supervisors place each bank against the thresholds of Appendix H.2 and Section 5.3. The same data would let carriers define and measure the slices of the accumulation they hold and trade, which diversified capital could then carry (Proposition 2).

One scenario across sectors. A provider failure hits insurers if cover is silent, firms and their lenders if cover is excluded, and providers to the extent that liability moves upstream. A single scenario, applied by insurance and banking supervisors under the contract wordings in force, would reveal where the loss falls, and the European Systemic Risk Board's work on systemic cyber risk offers a template (European Systemic Risk Board 2020). Because the factor moves from insurers to lenders as exclusions spread (Corollary 1), a scenario run only

Table 7. Terms of firms' relationships with foundation-model providers in SEC filings

Panel A. What the 114 firms that disclose dependence say in their annual reports

Firms whose annual reports	Firms
Name a provider or third-party models as a risk	65
as a dependence	27
naming outages or degraded service	23
naming discontinuation or loss of access	19
naming changes in price or terms	26
Describe an architecture that uses several models	10
State that switching provider would be slow, costly or uncertain	4
Refer to provider indemnities, all as incomplete	3
Refer to provider limits of liability	1
Report service-level commitments from a provider	0
Report insurance for AI output or for a provider's failure	0
Firms that disclose dependence on a provider	114

Panel B. Agreements with providers filed as exhibits

Agreement	Terms bearing on who bears a provider's failure
CoreWeave–OpenAI, services, 2025	Reserved-capacity fees; termination only for uncured breach or insolvency; on CoreWeave's insolvency OpenAI may move the contract to a substitute operator; no credit support from OpenAI
Cerebras–OpenAI, capacity, 2025	OpenAI commits to 250 MW by end-2026 and 500 MW by end-2027; delivery delays shorten the term; service-level agreement
Nscale–Anthropic, GPU services, 2026	Supplier finances chips and facility on terms set by Anthropic's credit profile, with no new liability for Anthropic; fees net of service credits; supplier indemnifies and insures
Akamai–Anthropic, services, 2026	Service credits are the customer's sole remedy for interruptions
Boost Run–Cohere, services, 2026	Committed-capacity fees invoiced monthly in advance; prepayments refunded if too little capacity is accepted by the outside date
TeraWulf landlord–Anthropic–Google, credit support, 2026	Google may, but need not, cure Anthropic's lease defaults or assume the lease
CoreWeave SPV, term loan, 2025	Revenue is the OpenAI contract; cash trap if OpenAI breaches or misses three months' payments; parent guarantee falls away once OpenAI is rated investment grade
CoreWeave SPV, term loan amendment, 2025	Coverage ratio of at least 1.40 on contracted cash flows from Microsoft and OpenAI
CoreWeave, revolving credit amendment, 2025	Debt secured on the OpenAI contract up to 75% loan-to-value, 90% if OpenAI is rated investment grade, at most USD 7 billion
Anthropic, secured revolving credit, 2026	Syndicated; 17 arranging banks, five of them European

Notes: Panel A: 228 annual reports filed from 2023 to 2026 by the 114 firms that state that their products run on a provider's models or that they use one internally, read in full; counts of firms. The first five rows count firms with a risk-factor paragraph that names a provider or third-party models and, for the indented rows, that uses the words listed; the last six rows are checked by hand against every paragraph on AI. Panel B: every agreement filed as an exhibit to an SEC filing from 2021 to November 2026 that names OpenAI, Anthropic, Cohere, Mistral or xAI by one of the legal names searched and sets terms between a provider and a supplier or lender, excluding equity, warrant, merger, governance and compensation agreements and those in which a provider appears only as a shareholder, investee or group affiliate (CoreWeave, Inc. 2025d; Cerebras Systems Inc. 2026; Nscale Ltd 2026; Akamai Technologies, Inc. 2026; Boost Run Inc. 2026; TeraWulf Inc. 2026; CoreWeave, Inc. 2025b,c,a; Anthropic, PBC 2026). Amounts redacted in the filings are not shown.

for insurers understates the exposure. For banks, the scenario reveals more than the capital formula does. The exposure is a concentration risk for lenders whose borrowers rely on the same provider, which supervisors could address under Pillar 2 once dependence is recorded, and the scenario should include the lending to AI infrastructure that is serviced from contracts with the same providers (Section 5.3). For insurers, own risk and solvency assessments could report the accumulation both stand-alone and within the book, under expected shortfall rather than value-at-risk, which can rank concentrated markets as safer (Lemma 3).

Exclusions, event limits and upstream capacity. The share of liability policies that carry AI exclusions measures how much of the loss distribution has left the insurance sector for firms and their lenders, which makes it a statistic for banking supervisors too. The model gives no case for mandatory cover once event limits and dependence data exist: limits per provider event keep most of the expected loss insured while capping the accumulation (Appendix C.3), and supervisors can help by recognising them, and reinsurance of the accumulation, in capital assessments. Upstream, a financial responsibility requirement for providers that scales with market share, placed among the AI Act's obligations for general-purpose models with systemic risk, would make recoveries under the Directive complete and independent of concentration (Proposition 4).

Predictions. Once dependence is recorded, the explanations of the exclusions in Section 2.4 make different predictions across carriers, which Appendix H.6 sets out.

8 Conclusion

Foundation-model concentration ties the liability losses of many firms to the fortunes of a few providers. When insurers exclude AI, that tie passes to the firms and, through their defaults, to their lenders, and Pillar 1 capital does not register how concentrated a lender's exposure is. Banks hold the same factor a second time as lenders to the infrastructure those providers contract for, and a failure severe enough to impair a provider's revenue reaches both sides in the same state, so a bank exposed to both takes both losses at once. In the European bank data the liability channel alone is small for the typical bank and for the euro-area system, and material only for lenders whose borrowers depend on the provider that fails; the combined exposure is larger, and at a loss rate of an eighth on infrastructure lending the asset channel matches the whole liability channel once such lending reaches 1.7% of a corporate book.

Neither channel is recorded where it would matter. The internal ratings-based formula sees the loss only once it reaches default rates, the Digital Operational Resilience Act reaches providers to financial institutions but not to their borrowers, and credit files do not record which model a borrower depends on. Public filings and loan data recover part of the picture. Eighty-nine firms state in annual reports that their products run on a provider's models, 73 of them on OpenAI's; 22 of them and four AI-infrastructure borrowers owe USD 78 billion in syndicated loans, held to a large extent by the same lenders on both channels; and when a provider has an incident, its clients' shares fall by about one and a half percentage points over two days relative to the clients of its competitors. Lenders charge these firms no premium in their loan spreads, their bonds do not move, and the firms diversify slowly, by adding

providers while keeping the first. Their contracts give them little recourse against a provider and their annual reports mention no cover for its failure, while the lenders to suppliers of computing capacity hold the providers' credit under covenants tied to their payments and ratings. Equity investors therefore register the provider factor, while loan and bond prices and the instruments of credit risk do not. The measured holdings are floors, because most adopters name no provider.

The quantities are stress values, because the probability of a liability-relevant failure and the share of a provider's clients it would harm have never been measured, and the loss rate on infrastructure exposure is a parameter rather than an estimate. The insurance-market evidence does not discriminate between explanations of the exclusions either, since the wordings observed so far are also what insurers clarifying silent cover would write. Records of which provider a firm depends on, and of which provider's contracts service the infrastructure a bank has lent against, collected at underwriting and in credit files, would not measure these probabilities, but they would measure each lender's exposure on both sides and allow the test the model proposes, that exclusions tighten with a carrier's exposure to the largest provider relative to its capital.

References

Acerbi, C. and D. Tasche (2002). On the coherence of expected shortfall. *Journal of Banking & Finance* 26(7), 1487–1503. doi:10.1016/S0378-4266(02)00283-2.

Akamai Technologies, Inc. (2026). *Master services agreement with Anthropic PBC, Exhibit 10.1 to Form 10-Q*. US Securities and Exchange Commission, 6 November 2026. https://www.sec.gov/Archives/edgar/data/1086222/000108622226000100/akam10q9302026ex101.htm.

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach (2022). The drivers of cyber risk. *Journal of Financial Stability* 60, 100989. doi:10.1016/j.jfs.2022.100989.

Aldasoro, I., L. Gambacorta, A. Korinek, V. Shreeti, and M. Stein (2024). *Intelligent financial system: How AI is transforming finance*. BIS Working Papers 1194. Bank for International Settlements.

Amazon Web Services (2025). *Summary of the Amazon DynamoDB service disruption in the Northern Virginia (US-EAST-1) Region*. Post-event summary, October 2025. https://aws.amazon.com/message/101925.

Amazon Web Services (2026a). *Amazon Bedrock FAQs*. Accessed 28 October 2026. https://aws.amazon.com/bedrock/faqs/.

Amazon Web Services (2026b). *AWS customer agreement*. Last updated 14 August 2026. https://aws.amazon.com/agreement/.

Andrade, G. and S. N. Kaplan (1998). How costly is financial (not economic) distress? Evidence from highly leveraged transactions that became distressed. *Journal of Finance* 53(5), 1443–1493. doi:10.1111/0022-1082.00062.

Anthropic (2025a). *A postmortem of three recent issues*. Engineering blog, 17 September 2025. https://www.anthropic.com/engineering/a-postmortem-of-three-recent-issues.

Anthropic (2025b). *Commercial terms of service*. Effective 17 June 2025. https://www.anthropic.com/legal/commercial-terms.

Anthropic (2026a). *Redeploying Fable 5*. 30 June 2026, updated 1 July 2026. https://www.anthropic.com/news/redeploying-fable-5.

Anthropic (2026b). *Statement on the US government directive to suspend access to Fable 5 and Mythos 5*. 12 June 2026. https://www.anthropic.com/news/fable-mythos-access.

Anthropic, PBC (2026). *Amended and restated credit agreement, Exhibit 10.10 to Form S-1*. US Securities and Exchange Commission, 26 October 2026. https://www.sec.gov/Archives/edgar/data/2130357/000162828026067168/exhibit1010-sx1.htm.

Armilla AI (2026). *Armilla AI raises Lloyd's-backed coverage to $25M as traditional insurers retreat from AI risk*. 21 January 2026. https://www.armilla.ai/resources/armilla-ai-raises-lloyds-backed-coverage-to-25m-as-traditional-insurers-retreat-from-ai-risk.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance* 9(3), 203–228. doi:10.1111/1467-9965.00068.

Aslan, I. (2026a). Capital destruction and the limits of insurability: What prices reveal about catastrophic risk from artificial intelligence. Working paper. doi:10.13140/RG.2.2.25664.70402.

Aslan, I. (2026b). *Correlated AI losses: Code and data*. Zenodo. Concept DOI, resolves to the latest version. doi:10.5281/zenodo.23733783.

Aslan, I. (2026c). From peril to security: A framework for assessing ILS readiness, with application to systemic cyber risk. SSRN Working Paper 6684118.

Aslan, I. (2026d). The insurability frontier: Ambiguity, capital, and the supply of insurance. Working paper.

Bank of England (2025). *Financial Stability in Focus: Artificial intelligence in the financial system*. Report. Bank of England. https://www.bankofengland.co.uk/financial-stability-in-focus/2025/april-2025.

Bank of England (2026). *Financial stability report*. July 2026. https://www.bankofengland.co.uk/financial-stability-report/2026/july-2026.

Barrot, J.-N. and J. Sauvagnat (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics* 131(3), 1543–1592. doi:10.1093/qje/qjw018.

Bartz v. Anthropic (2026). *Order granting final approval of class action settlement*. No. 3:24-cv-05417 (N.D. Cal.), 20 July 2026.

Basel Committee on Banking Supervision (2005). *An Explanatory Note on the Basel II IRB Risk Weight Functions*. Report. Bank for International Settlements. https://www.bis.org/bcbs/irbriskweight.pdf.

Battiston, S., A. Mandel, I. Monasterolo, F. Schütze, and G. Visentin (2017). A climate stress-test of the

financial system. *Nature Climate Change* 7(4), 283–288.

Berger, A. N., F. Curti, A. Mihov, and J. Sedunov (2022). Operational risk is more systemic than you think: Evidence from U.S. bank holding companies. *Journal of Banking & Finance* 143, 106619. doi:10.1016/j.jbankfin.2022.106619.

Berliner, B. (1982). *Limits of Insurability of Risks*. Englewood Cliffs, NJ, Prentice-Hall.

Bianchi, J. (2011). Overborrowing and systemic externalities in the business cycle. *American Economic Review* 101(7), 3400–3426.

Biener, C., M. Eling, and J. H. Wirfs (2015). Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance – Issues and Practice* 40(1), 131–158. doi:10.1057/gpp.2014.19.

Bloomberg Law (2026). *Insurer AI exclusions spark policyholder alarm on coverage gaps*. 6 July 2026. https://news.bloomberglaw.com/insurance/insurer-ai-exclusions-spark-policyholder-alarm-on-coverage-gaps.

Board of Governors of the Federal Reserve System (2023). *Review of the Federal Reserve's supervision and regulation of Silicon Valley Bank*. Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/publications/files/svb-review-20230428.pdf.

Bommasani, R., K. A. Creel, A. Kumar, D. Jurafsky, and P. Liang (2022). Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? In *Advances in Neural Information Processing Systems*, pp. 3663–3678. https://proceedings.neurips.cc/paper_files/paper/2022/hash/17a234c91f746d9625a75cf8a8731ee2-Abstract-Conference.html.

Boost Run Inc. (2026). *Service agreement with Cohere Inc., Exhibit 10.1 to Form 8-K*. US Securities and Exchange Commission, 6 October 2026. https://www.sec.gov/Archives/edgar/data/2090646/000149315226045861/ex10-1.htm.

Boyer, M. and J.-J. Laffont (1997). Environmental risks and bank liability. *European Economic Review* 41(8), 1427–1459. doi:10.1016/S0014-2921(96)00034-7.

Buch, C. (2026). *Letter to Fabio De Masi, Member of the European Parliament, on the financial stability risks of advanced AI applications (QZ-010)*. ECB Banking Supervision, 25 June 2026. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.mepletter260625_DeMasi~b3f7e17e42.en.pdf.

Buiten, M. C., A. de Streel, and M. Peitz (2023). The law and economics of AI liability. *Computer Law & Security Review* 48, 105794. doi:10.1016/j.clsr.2023.105794.

Business Insurance (2025). *OpenAI, Anthropic eye investor funds to settle AI lawsuits: Report*. October 2025. https://www.businessinsurance.com/openai-anthropic-eye-investor-funds-to-settle-ai-lawsuits-report/.

Carroll, S. J., D. Hensler, J. Gross, E. M. Sloss, M. Schonlau, A. Abrahamse, and J. S. Ashwood (2005). *Asbestos Litigation*. Santa Monica, CA, RAND Corporation.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics* 136(2), 1255–1321.

Cerebras Systems Inc. (2026). *Master relationship agreement with OpenAI OpCo, LLC, Exhibit 10.11 to Form S-1/A*. US Securities and Exchange Commission, 4 May 2026. https://www.sec.gov/Archives/edgar/data/2021728/000162828026029503/exhibit1011-sx1a.htm.

Cifuentes, A. and G. O'Connor (1996). *The binomial expansion method applied to CBO/CLO analysis*. Special report. Moody's Investors Service.

Cloud Security Alliance (2026). *Foundation model concentration: The uninsurable AI risk*. CSA research note, AI Safety Initiative, 6 July 2026. https://labs.cloudsecurityalliance.org/research/csa-research-note-foundation-model-concentration-insurance-r/.

Cohere (2026). *Software as a service agreement*. Undated web page. https://cohere.com/saas-agreement.

CoreWeave, Inc. (2025a). *Amendment no. 4 to revolving credit agreement, Exhibit 10.7 to Form 10-Q*. US Securities and Exchange Commission, 13 November 2025. https://www.sec.gov/Archives/edgar/data/1769628/000176962825000062/exhibit107.htm.

CoreWeave, Inc. (2025b). *Credit agreement of CoreWeave Compute Acquisition Co. VII, LLC, Exhibit 10.1 to Form 8-K*. US Securities and Exchange Commission, 31 July 2025. https://www.sec.gov/Archives/edgar/data/1769628/000176962825000033/coreweave30-ex101xtermlo.htm.

CoreWeave, Inc. (2025c). *Fifth amendment to credit agreement of CoreWeave Compute Acquisition Co. IV, LLC, Exhibit 10.1 to Form 8-K*. US Securities and Exchange Commission, 2 October 2025. https://www.sec.gov/Archives/edgar/data/1769628/000119312525227562/d910811dex101.htm.

CoreWeave, Inc. (2025d). *Master services agreement with OpenAI OpCo, LLC, Exhibit 10.24 to Form S-1/A*. US Securities and Exchange Commission, 12 March 2025. https://www.sec.gov/Archives/edgar/d

ata/1769628/000119312525052207/d899798dex1024.htm.

Cortés, K. R. and P. E. Strahan (2017). Tracing out capital flows: How financially integrated banks respond to natural disasters. *Journal of Financial Economics* 125(1), 182–199.

Council of the European Communities (1985). *Council Directive 85/374/EEC on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products*. Official Journal of the European Communities L 210, 7 August 1985.

Crosignani, M., M. Macchiavelli, and A. F. Silva (2023). Pirates without borders: The propagation of cyberattacks through firms' supply chains. *Journal of Financial Economics* 147(2), 432–448.

CrowdStrike (2026). *Terms and conditions*. Web page, version of 7 January 2026. https://www.crowdstrike.com/terms-conditions/.

CrowdStrike Holdings (2025). *Fourth quarter fiscal 2025 earnings call*. Transcript, 4 March 2025. https://www.nasdaq.com/articles/crowdstrike-crwd-q4-2025-earnings-call-transcript.

CrowdStrike Holdings (2026). *Form 10-Q for the quarter ended 31 July 2026: commitments and contingencies*. US Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1535527/000153552726000031/R18.htm.

CrowdStrike v. Delta (2024). *Complaint for declaratory judgment*. No. 1:24-cv-04904 (N.D. Ga.), 25 October 2024. https://storage.courtlistener.com/recap/gov.uscourts.gand.335602/gov.uscourts.gand.335602.1.0.pdf.

CyberCube (2024). *CyberCube estimates $400mn–$1.5bn global insured losses from the CrowdOut event*. Press release, Business Wire, 25 July 2024. https://www.businesswire.com/news/home/20240725566739/en.

CyberCube (2025). *CyberCube reveals insurance loss estimate for AWS "Amazonk" outage*. Press release, Business Wire, 23 October 2025. https://www.businesswire.com/news/home/20251023532923/en/CyberCube-Reveals-Insurance-Loss-Estimate-for-AWS-Amazonk-Outage.

Daníelsson, J., R. Macrae, and A. Uthemann (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance* 140, 106290. doi:10.1016/j.jbankfin.2021.106290.

Daníelsson, J., B. N. Jorgensen, G. Samorodnitsky, M. Sarma, and C. G. de Vries (2013). Fat tails, VaR and subadditivity. *Journal of Econometrics* 172(2), 283–291.

DatacenterDynamics (2025). *CoreWeave closes $2.6bn in secured debt financing facility*. 31 July 2025. https://www.datacenterdynamics.com/en/news/coreweave-closes-26bn-in-secured-debt-financing-facility/.

Davydenko, S. A., I. A. Strebulaev, and X. Zhao (2012). A market-based study of the cost of default. *Review of Financial Studies* 25(10), 2959–2999. doi:10.1093/rfs/hhs091.

Del Rio v. CrowdStrike (2026). *Judgment of the Court of Appeals for the Fifth Circuit, per curiam*. No. 25-50518, 20 May 2026. https://www.ca5.uscourts.gov/opinions/unpub/25/25-50518.0.pdf.

Delta Air Lines (2024). *Form 8-K, 8 August 2024*. US Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/27904/000168316824005369/delta_8k.htm.

Deutscher Bundestag (2026). *Entwurf eines Gesetzes zur Modernisierung des Produkthaftungsrechts und zur Änderung weiterer Vorschriften*. Drucksache 21/4297, 25 February 2026; adopted 8 October 2026. https://dserver.bundestag.de/btd/21/042/2104297.pdf.

Dhaene, J., S. Vanduffel, M. J. Goovaerts, R. Kaas, Q. Tang, and D. Vyncke (2006). Risk measures and comonotonicity: A review. *Stochastic Models* 22(4), 573–606.

Duarte, E., J. Zuloaga, C. Cohen, and R. Martin (2026). *ECB steps up scrutiny of European banks' AI industry exposure*. Bloomberg, 23 February 2026, republished by Insurance Journal. https://www.insurancejournal.com/news/international/2026/02/23/859222.htm.

Duffie, D. and J. Younger (2019). *Cyber runs: How a cyber attack could affect U.S. financial institutions*. Hutchins Center Working Paper 51. Brookings Institution. https://www.brookings.edu/articles/cyber-runs/.

Düllmann, K. and N. Masschelein (2007). A tractable model to measure sector concentration risk in credit portfolios. *Journal of Financial Services Research* 32(1), 55–79.

Eisenbach, T. M., A. Kovner, and M. J. Lee (2022). Cyber risk and the U.S. financial system: A pre-mortem analysis. *Journal of Financial Economics* 145(3), 802–826.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang (2023). *Generative AI and firm values*. NBER Working Paper 31222. National Bureau of Economic Research. https://www.nber.org/papers/w31222.

Eling, M. and J. Wirfs (2019). What are the actual costs of cyber risk events? *European Journal of Operational Research* 272(3), 1109–1119. doi:10.1016/j.ejor.2018.07.021.

European Banking Authority (2025). *2025 EU-wide transparency exercise*. Full database, 119 banks, reference dates 30 September 2024 to 30 June 2025. https://www.eba.europa.eu/risk-and-data-analysis/risk-analysis/eu-wide-transparency-exercise.

European Central Bank (2025). *Aggregated results of the 2025 Supervisory Review and Evaluation Process*. ECB Banking Supervision, 18 November 2025. https://www.bankingsupervision.europa.eu/activities/srep/2025/html/ssm.srep202511_aggregatedresults2025.en.html.

European Central Bank (2026a). *Financial Stability Review, May 2026*. European Central Bank. https://www.ecb.europa.eu/press/financial-stability-publications/fsr/html/ecb.fsr202605~50566915a7.en.html.

European Central Bank (2026b). *Supervisory banking statistics: Second quarter 2026*. European Central Bank, Banking Supervision. 15 September 2026. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.supervisorybankingstatistics_second_quarter_2026_202609.en.pdf.

European Central Bank (2026c). *The euro area bank lending survey: Third quarter of 2026*. European Central Bank. 27 October 2026. https://www.ecb.europa.eu/stats/ecb_surveys/bank_lending_survey/html/ecb.blssurvey2026q3~c6a731e213.en.html.

European Commission (2015). *Commission Delegated Regulation (EU) 2015/35 supplementing Directive 2009/138/EC*. Official Journal of the European Union L 12, 17 January 2015, as amended.

European Commission (2025). *Withdrawal of Commission proposals*. Official Journal of the European Union C/2025/5423, 6 October 2025. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:C_202505423.

European Parliament (2020). *Resolution of 20 October 2020 with recommendations to the Commission on a civil liability regime for artificial intelligence*. 2020/2014(INL), P9_TA(2020)0276. https://www.europarl.europa.eu/doceo/document/TA-9-2020-0276_EN.html.

European Parliament and Council (2004). *Directive 2004/35/CE on environmental liability with regard to the prevention and remedying of environmental damage*. Official Journal of the European Union L 143, 30 April 2004.

European Parliament and Council (2009). *Directive 2009/138/EC on the taking-up and pursuit of the business of insurance and reinsurance (Solvency II)*. Official Journal of the European Union L 335, 17 December 2009.

European Parliament and Council (2024). *Directive (EU) 2024/2853 on liability for defective products and repealing Council Directive 85/374/EEC*. Official Journal of the European Union, L series, 18 November 2024. https://eur-lex.europa.eu/eli/dir/2024/2853/oj/eng.

European Parliament and Council (2026). *Regulation (EU) 2026/1744 amending Regulation (EU) 2024/1689 as regards the simplification of the implementation of harmonised rules on artificial intelligence*. Official Journal of the European Union, L series, 24 July 2026. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32026R1744.

European Supervisory Authorities (2025). *European Supervisory Authorities designate critical ICT third-party providers under the Digital Operational Resilience Act*. 18 November 2025. https://www.eiopa.europa.eu/european-supervisory-authorities-designate-critical-ict-third-party-providers-under-digital-2025-11-18_en.

European Systemic Risk Board (2020). *Systemic cyber risk*. Report. European Systemic Risk Board. https://www.esrb.europa.eu/pub/pdf/reports/esrb.report200219_systemiccyberrisk~101a09685e.en.pdf.

European Systemic Risk Board (2025). *Artificial intelligence and systemic risk*. Report. Advisory Scientific Committee, European Systemic Risk Board. https://www.esrb.europa.eu/news/pr/date/2025/html/esrb.pr251204~d8ef41c2f1.en.html.

Eurostat (2025). *Use of artificial intelligence in enterprises*. Statistics Explained, data extracted December 2025. https://ec.europa.eu/eurostat/statistics-explained/index.php/Use_of_artificial_intelligence_in_enterprises.

Eurostat (2026a). *Artificial intelligence by size class of enterprise*. Online data code isoc_eb_ai, updated 15 June 2026. https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ai/default/table.

Eurostat (2026b). *Artificial intelligence by NACE Rev. 2 activity*. Online data code isoc_eb_ain2. https://ec.europa.eu/eurostat/databrowser/view/isoc_eb_ain2/default/table.

Federal Republic of Germany (1989). *Gesetz über die Haftung für fehlerhafte Produkte (Produkthaftungsgesetz)*. BGBl. I S. 2198, as amended. https://www.gesetze-im-internet.de/prodhaftg/__10.html.

Federal Republic of Germany (2026). *Gesetz zur Modernisierung des Produkthaftungsrechts und zur*

Änderung weiterer Vorschriften. 22 October 2026, BGBl. 2026 I Nr. 321, 27 October 2026. https://www.recht.bund.de/bgbl/1/2026/321/regelungstext.pdf.

Financial Stability Board (2019). *Third-party dependencies in cloud services: Considerations on financial stability implications*. 9 December 2019. https://www.fsb.org/uploads/P091219-2.pdf.

Financial Stability Board (2023). *Enhancing Third-Party Risk Management and Oversight: A toolkit for financial institutions and financial authorities*. Final report. Financial Stability Board. https://www.fsb.org/uploads/P041223-1.pdf.

Financial Stability Board (2024). *The Financial Stability Implications of Artificial Intelligence*. Report. Financial Stability Board. https://www.fsb.org/uploads/P14112024.pdf.

FINRA (2026). *TRACE Enhanced*. Corporate bond transactions database, accessed through Wharton Research Data Services on 18 November 2026. https://wrds-www.wharton.upenn.edu.

Florackis, C., C. Louca, R. Michaely, and M. Weber (2023). Cybersecurity risk. *Review of Financial Studies* 36(1), 351–407. doi:10.1093/rfs/hhac024.

Froot, K. A. (2001). The market for catastrophe risk: a clinical examination. *Journal of Financial Economics* 60(2–3), 529–571. doi:10.1016/S0304-405X(01)00052-6.

Froot, K. A., D. S. Scharfstein, and J. C. Stein (1993). Risk management: Coordinating corporate investment and financing policies. *Journal of Finance* 48(5), 1629–1658.

Froot, K. A. and J. C. Stein (1998). Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics* 47(1), 55–82. doi:10.1016/S0304-405X(97)00037-8.

Gallagher Re (2026). *Smart systems, blind spots: Rethinking insurance for the AI era*. Gallagher Re. March 2026. https://www.ajg.com/gallagherre/-/media/files/gallagher/gallagherre/news-and-insights/2026/march/rethinking-insurance-for-the-ai-era.pdf.

Gallin, L. (2026). *New highs for traditional and alternative takes global reinsurer capital to record $785bn: Aon*. Reinsurance News, 1 April 2026. https://www.reinsurancene.ws/new-highs-for-traditional-alternative-takes-global-reinsurer-capital-to-record-785bn-aon/.

Gambacorta, L. and H. S. Shin (2018). Why bank capital matters for monetary policy. *Journal of Financial Intermediation* 35, 17–29.

Garmaise, M. J. and T. J. Moskowitz (2009). Catastrophic risk and credit markets. *Journal of Finance* 64(2), 657–707.

Gensler, G. and L. Bailey (2020). Deep learning and financial stability. Working paper, SSRN 3723132. doi:10.2139/ssrn.3723132.

Germond, N. (2026). *Discover the endorsement and exclusions filing trends for 2026*. Independent Insurance Agents & Brokers of America, 22 July 2026, reporting filing data from Insuraviews. https://www.independentagent.com/?p=51022.

Google Cloud (2026). *Google Cloud Platform terms of service and service specific terms*. Terms of service last modified 12 October 2026. https://cloud.google.com/terms.

Gordy, M. B. (2003). A risk-factor model foundation for ratings-based bank capital rules. *Journal of Financial Intermediation* 12(3), 199–232. doi:10.1016/S1042-9573(03)00040-8.

Gordy, M. B. and E. Lütkebohmert (2013). Granularity adjustment for regulatory capital assessment. *International Journal of Central Banking* 9(3), 38–77.

Gouriéroux, C., J.-P. Laurent, and O. Scaillet (2000). Sensitivity analysis of values at risk. *Journal of Empirical Finance* 7(3–4), 225–245. doi:10.1016/S0927-5398(00)00011-6.

Guy Carpenter (2024). *A global outage with widespread impact*. Guy Carpenter Insights, 19 July 2024, updated 1 August 2024. https://www.guycarp.com/insights/2024/07/global-outage-with-widespread-impact.html.

Hacker, P. (2023). The European AI liability directives – Critique of a half-hearted approach and lessons for the future. *Computer Law & Security Review* 51, 105871. doi:10.1016/j.clsr.2023.105871.

Harrington, J. S. (2026). *AI exclusions*. Rough Notes, 22 January 2026. https://roughnotes.com/ai-exclusions/.

Ibragimov, R., D. Jaffee, and J. Walden (2009). Nondiversification traps in catastrophe insurance markets. *Review of Financial Studies* 22(3), 959–993. doi:10.1093/rfs/hhn021.

Ibragimov, R., D. Jaffee, and J. Walden (2011). Diversification disasters. *Journal of Financial Economics* 99(2), 333–348. doi:10.1016/j.jfineco.2010.08.015.

Independent Insurance Agents & Brokers of America (2025). *Verisk to roll out new general liability exclusions for generative AI exposures*. 21 October 2025. https://www.independentagent.com/vu_reso

urce/verisk-to-roll-out-new-general-liability-exclusions-for-generative-ai-exposures/.

Insurance Edge (2025). *Parametrix estimates AWS outage costs*. 30 October 2025. https://insurance-edge.net/2025/10/30/parametrix-estimates-aws-outage-costs/.

IREN Limited (2026). *Current report on Form 8-K: IE US Hardware 3 financing*. US Securities and Exchange Commission, 29 May 2026. https://www.sec.gov/Archives/edgar/data/1878848/000114036126023427/ef20075181_8k.htm.

Jaffee, D. M. and T. Russell (1997). Catastrophe insurance, capital markets, and uninsurable risks. *Journal of Risk and Insurance* 64(2), 205–230. doi:10.2307/253729.

Jamilov, R., H. Rey, and A. Tahoun (2021). *The anatomy of cyber risk*. NBER Working Paper 28906. National Bureau of Economic Research. https://www.nber.org/papers/w28906.

Kalkbrener, M. (2005). An axiomatic approach to capital allocation. *Mathematical Finance* 15(3), 425–437.

Kashyap, A. K. and A. Wetherilt (2019). Some principles for regulating cyber risk. *AEA Papers and Proceedings* 109, 482–487. doi:10.1257/pandp.20191058.

Kleinberg, J. and M. Raghavan (2021). Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118(22), e2018340118. doi:10.1073/pnas.2018340118.

Koetter, M., F. Noth, and O. Rehbein (2020). Borrowers under water! Rare disasters, regional banks, and recovery lending. *Journal of Financial Intermediation* 43.

Leung, A., R. Zhang, E. Ling, K. Toyoda, and S. Loh (2026). *The insurability frontier of AI risk: Mapping threats to affirmative coverage, silent exposures, and exclusions*. arXiv preprint arXiv:2605.18784. https://arxiv.org/abs/2605.18784.

Leytes, A. (2026). *Cyber-financial contagion: Modeling the propagation of an AI vendor compromise through the banking system*. arXiv:2609.10350. https://arxiv.org/abs/2609.10350.

Lior, A. (2022). Insuring AI: The role of insurance in artificial intelligence regulation. *Harvard Journal of Law & Technology* 35(2), 467. https://jolt.law.harvard.edu/assets/articlePDFs/v35/2.-Lior-Insuring-AI.pdf.

Lloyd's (2019). *Market Bulletin Y5258: Providing clarity for Lloyd's customers on coverage for cyber exposures*. 4 July 2019.

Lloyd's (2020). *Market Bulletin Y5277: Update – providing clarity for Lloyd's customers on coverage for cyber exposures*. 29 January 2020.

Lloyd's (2022). *Market Bulletin Y5381: State backed cyber-attack exclusions*. 16 August 2022.

Lloyd's (2026). *RDS 2026: Realistic disaster scenarios, scenario specification*. Document EM 568, version 1.0, January 2026. https://www.lloyds.com/market-resources/underwriting/realistic-disaster-scenarios-rds.

Lloyd's and AIR Worldwide (2018). *Cloud Down: Impacts on the US economy*. Emerging Risk Report 2018 – Technology. Lloyd's. https://assets.lloyds.com/media/4f1c0e52-0263-43e3-860d-b1cc10f080ee/pdf-air-cyber-lloyds-public-2018-final.pdf.

Lockton Re and Armilla AI (2026). *Ready or not: The impact of artificial intelligence on insurance risks*. Lockton Re. 5 February 2026. https://global.lockton.com/re/en/news-insights/new-lockton-re-and-armilla-report-ready-or-not-finds-ai-is-reshaping.

Loizos, C. (2025). *AI is too risky to insure, say people whose job is insuring risk*. TechCrunch, 23 November 2025, reporting the Financial Times. https://techcrunch.com/2025/11/23/ai-is-too-risky-to-insure-say-people-whose-job-is-insuring-risk/.

Lopez, O. and D. Nkameni (2026). *Cloud failure and cyber insurance: Calibration of stress scenarios and diversification*. arXiv:2607.18815. https://arxiv.org/abs/2607.18815.

Lorenzoni, G. (2008). Inefficient credit booms. *Review of Economic Studies* 75(3), 809–833.

LSEG (2026). *DealScan*. Syndicated loan database, accessed through Wharton Research Data Services on 11 November 2026. https://wrds-www.wharton.upenn.edu.

Marshall, A. W., I. Olkin, and B. C. Arnold (2011). *Inequalities: Theory of Majorization and Its Applications*, Second edition. New York, NY, Springer. doi:10.1007/978-0-387-68276-1.

Maschke, M. C., S. D. Nelson, and J. W. Simek (2026). *Artificial intelligence exclusions are creeping into insurance, but cyber policies aren't the issue (yet)*. Big "I" Virtual University, Independent Insurance Agents & Brokers of America, undated, retrieved 28 October 2026. https://www.independentagent.com/?p=45658.

Mayers, D. and C. W. Smith (1982). On the corporate demand for insurance. *Journal of Business* 55(2), 281–296.

Meta (2025). *Llama 4 community license agreement*. Effective 5 April 2025. https://www.llama.com/llam

a4/license/.

Meyers, G. and P. Shi (2011). *Loss reserving data pulled from NAIC Schedule P*. Casualty Actuarial Society, accident years 1988 to 1997. https://www.casact.org/publications-research/research/research-resources/loss-reserving-data-pulled-naic-schedule-p.

Microsoft (2019). *Microsoft online subscription agreement*. Last updated March 2019. https://azure.microsoft.com/en-us/support/legal/subscription-agreement/.

Microsoft (2023). *A statement from Microsoft Chairman and CEO Satya Nadella*. Official Microsoft Blog, posted 17 November 2023, with updates of 19 and 21 November 2023. https://blogs.microsoft.com/blog/2023/11/17/a-statement-from-microsoft-chairman-and-ceo-satya-nadella/.

Microsoft (2026). *Product terms: universal license terms for online services*. Version current on 28 October 2026. https://www.microsoft.com/licensing/terms/product/ForOnlineServices/all.

Mistral AI (2026). *Commercial terms of service*. Effective 25 September 2026. https://legal.mistral.ai/terms/commercial-terms-of-service.

Mobley v. Workday (2024). *Order granting in part and denying in part motion to dismiss*. No. 3:23-cv-00770-RFL (N.D. Cal.), 12 July 2024, 740 F. Supp. 3d 796.

Mobley v. Workday (2025). *Order granting preliminary collective certification*. No. 3:23-cv-00770-RFL (N.D. Cal.), 16 May 2025.

Moffatt v. Air Canada (2024). *Decision of the Civil Resolution Tribunal of British Columbia*. 2024 BCCRT 149, 14 February 2024. https://decisions.civilresolutionbc.ca/crt/crtd/en/525448/1/document.do.

National Association of Insurance Commissioners (2026a). *2025 Annual Property and Casualty and Title Insurance Industries Analysis Report*. National Association of Insurance Commissioners. https://content.naic.org/sites/default/files/2025-annual-property-and-casualty-and-title-insurance-industries-analysis-report.pdf.

National Association of Insurance Commissioners (2026b). *Property and casualty insurance industry: 2025 top 25 groups and companies by countrywide premium*. Market share report, based on filings received by 18 March 2026. https://content.naic.org/sites/default/files/research-actuarial-property-casualty-market-share.pdf.

Nscale Ltd (2026). *Order for GPU services with Anthropic PBC, Exhibit 10.25 to Form S-1*. US Securities and Exchange Commission, 18 September 2026. https://www.sec.gov/Archives/edgar/data/2110365/000119312526395475/ck0002110365-ex10_25.htm.

OpenAI (2023). *OpenAI announces leadership transition*. 17 November 2023. https://openai.com/index/openai-announces-leadership-transition/.

OpenAI (2024). *New embedding models and API updates*. 25 January 2024. https://openai.com/index/new-embedding-models-and-api-updates/.

OpenAI (2026a). *Service terms*. Version of 12 June 2026. https://openai.com/policies/service-terms/.

OpenAI (2026b). *Services agreement*. Updated 1 December 2025, effective 1 January 2026. https://openai.com/policies/services-agreement/.

Parametrix (2024). *CrowdStrike's impact on the Fortune 500*. Report and press release, 24 July 2024. https://parametrixinsurance.com/reports-white-papers/crowdstrikes-impact-on-the-fortune-500.

Pitchford, R. (1995). How liable should a lender be? The case of judgment-proof firms and environmental risk. *American Economic Review* 85(5), 1171–1186. https://ideas.repec.org/a/aea/aecrev/v85y1995i5p1171-86.html.

Prudential Regulation Authority (2017). *Cyber insurance underwriting risk*. Supervisory Statement SS4/17, 5 July 2017. https://www.bankofengland.co.uk/prudential-regulation/publication/2017/cyber-insurance-underwriting-risk-ss.

Pykhtin, M. (2004). Multi-factor adjustment. *Risk* 17, 85–90.

Ramp (2026). *Ramp AI Index*. Token-spend shares for the week ending 23 October 2026, retrieved on 28 October 2026. https://ramp.com/data/ai-index.

Reinsurance News (2026). *US life insurers increase AI investments as exposure risks grow, says Moody's*. 23 October 2026. https://www.reinsurancene.ws/us-life-insurers-increase-ai-investments-as-exposure-risks-grow-says-moodys/.

Rosanes, M. (2026). *The AI clause that may already be on your contractor client's policy*. Insurance Business, 25 August 2026, reporting an analysis of state filing records by Trades Coverage. https://www.insurancebusinessmag.com/us/news/sme/the-ai-clause-that-may-already-be-on-your-contractor-clients-policy-587395.aspx.

Rossa, R. and D. Read (2025). *EU regulators designate critical ICT third-party providers under DORA*. TLT

LLP, 4 December 2025. https://www.tlt.com/insights-and-events/insight/eu-regulators-designate-critical-ict-third-party-providers-under-dora.

Scaillet, O. (2004). Nonparametric estimation and sensitivity analysis of expected shortfall. *Mathematical Finance* 14(1), 115–129.

Semafor (2023). *Sam Altman is returning to OpenAI*. 21 November 2023, updated 22 November 2023. https://www.semafor.com/article/11/21/2023/sam-altman-is-returning-to-openai.

Shavell, S. (1986). The judgment proof problem. *International Review of Law and Economics* 6(1), 45–58. doi:10.1016/0144-8188(86)90038-4.

Stempel, J. (2026). *CrowdStrike defeats shareholder lawsuit over huge software outage*. Reuters, 14 January 2026. https://in.marketscreener.com/news/crowdstrike-defeats-shareholder-lawsuit-over-huge-software-outage-ce7e58d9d98bf12c.

Swiss Federal Council (2024). *Verordnung über die Beaufsichtigung von privaten Versicherungsunternehmen (Aufsichtsverordnung, AVO), SR 961.011*. Revised version in force since 1 January 2024. https://www.fedlex.admin.ch/eli/cc/2005/735/de.

Tasche, D. (2008). *Capital allocation to business units and sub-portfolios: The Euler principle*. arXiv:0708.2542. https://arxiv.org/abs/0708.2542.

Temkin, M. (2024). *Menlo Ventures and Anthropic team up on a $100M AI fund*. TechCrunch, 17 July 2024. https://techcrunch.com/2024/07/17/menlo-ventures-and-anthropic-team-up-on-a-100m-ai-fund.

TeraWulf Inc. (2026). *Form of credit support agreement among Raylan Data LLC, TeraWulf Inc., Anthropic, PBC and Google LLC, Exhibit 10.1 to Form 8-K*. US Securities and Exchange Commission, 4 November 2026. https://www.sec.gov/Archives/edgar/data/1083301/000162828026069810/exhibit101-googlexterawulf.htm.

Toups, C., R. Bommasani, K. A. Creel, S. H. Bana, D. Jurafsky, and P. Liang (2023). Ecosystem-level analysis of deployed machine learning reveals homogeneous outcomes. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*.

Trout, C., S. Koyejo, S. Romanosky, G. Ripamonti, L. Thompson, D. Spain, A. Taylor, K. Casey, S. Casper, and M. Botvinick (2026). *Underwriting the agent economy: The blueprint for an AI insurance stack*. arXiv:2607.11999. https://arxiv.org/abs/2607.11999.

Tully, T., J. Redfern, D. Das, and D. Xiao (2025a). *2025 mid-year LLM market update: Foundation model landscape and economics*. Menlo Ventures. 31 July 2025. https://menlovc.com/perspective/2025-mid-year-llm-market-update/.

Tully, T., J. Redfern, D. Das, and D. Xiao (2025b). *2025: The state of generative AI in the enterprise*. Menlo Ventures. 9 December 2025. https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/.

U.S. Department of Justice and Federal Trade Commission (2023). *Merger Guidelines*. 18 December 2023.

U.S. Securities and Exchange Commission (2026). *EDGAR full text search*. Form 10-K filings; queries run on 27 October 2026. https://www.sec.gov/edgar/search/.

Vasicek, O. (2002). Loan portfolio value. *Risk*, 160–162. https://www.risk.net/risk-management/credit-risk/1500333/loan-portfolio-value.

Wang, S., J. Kahl, and S. Xu (2026). *Leaders, gainers and unexpected winners in the enterprise AI arms race*. Andreessen Horowitz, 30 January 2026. https://www.a16z.news/p/leaders-gainers-and-unexpected-winners.

Welburn, J. W. and A. M. Strong (2022). Systemic cyber risk and aggregate impacts. *Risk Analysis* 42(8), 1606–1622. doi:10.1111/risa.13715.

Weston, D. (2024). *Helping our customers through the CrowdStrike outage*. Official Microsoft Blog, 20 July 2024. https://blogs.microsoft.com/blog/2024/07/20/helping-our-customers-through-the-crowdstrike-outage/.

Wiggers, K. (2025). *Meta previews an API for its Llama AI models*. TechCrunch, 29 April 2025. https://techcrunch.com/2025/04/29/meta-previews-an-api-for-its-llama-ai-models/.

Wilkinson, C. (2023). *Beazley planning to determine risk appetite for AI*. Business Insurance, 3 November 2023. https://www.businessinsurance.com/beazley-planning-to-determine-risk-appetite-for-ai/.

Yahoo Finance (2026). *Historical prices*. Daily closing and adjusted closing prices, retrieved on 28 October and 11 November 2026. https://finance.yahoo.com.

Zanjani, G. (2002). Pricing and capital allocation in catastrophe insurance. *Journal of Financial Economics* 65(2), 283–305. doi:10.1016/S0304-405X(02)00141-1.

Zhu, Q. (2026). *Insurance of agentic AI*. arXiv:2606.05449. https://arxiv.org/abs/2606.05449.

A Legal and prudential background, provider contracts and the CrowdStrike outage

This appendix gives the legal and prudential detail summarised in Section 2.5 and the evidence on contracts and on the CrowdStrike outage summarised in Section 2.3.

A.1 The revised Product Liability Directive

Directive (EU) 2024/2853 replaces the Product Liability Directive of 1985. It entered into force on 8 December 2024, must be transposed by 9 December 2026, and applies to products placed on the market or put into service after that date (European Parliament and Council 2024, Arts. 2(1) and 22). Software is a product however it is supplied, including through cloud technologies or as a service (Art. 4(1) and Recital 13), and developers of software, including providers of AI systems within the meaning of the AI Act, are to be treated as manufacturers (Recital 13); free and open-source software supplied outside a commercial activity is excluded (Art. 2(2)). Whether a general-purpose model reached through an interface is a product in its own right, or a component of the product into which a deployer integrates it, has not yet been tested in court. The right to compensation belongs to natural persons (Art. 5) and covers three kinds of damage: death and personal injury, including medically recognised damage to psychological health; damage to property not used exclusively for professional purposes; and the destruction or corruption of data not used for professional purposes (Art. 6). Pure economic loss is not covered. The manufacturer of a defective component integrated into, or inter-connected with, a product within the manufacturer's control is liable alongside the manufacturer of the product, and the two are jointly and severally liable (Arts. 8(1) and 12(1)), with recourse between them governed by national law (Art. 14). A defect introduced by a software update within the manufacturer's control does not escape liability (Art. 11(2)). Courts can order disclosure of evidence (Art. 9), and defectiveness or causation is presumed where technical or scientific complexity makes proof excessively difficult and the claimant shows that defectiveness or causation is likely (Art. 10(4)), a provision the recitals connect explicitly to machine learning (Recital 48); the presumptions are rebuttable (Art. 10(5)).

Liability cannot be limited or excluded, by contract or by national law, in relation to the injured person (Art. 15), and the old option to cap it has gone. The 1985 Directive allowed Member States to limit a manufacturer's total liability for death and personal injury caused by identical items with the same defect to no less than ECU 70 million (Council of the European Communities 1985, Art. 16); Germany set the ceiling at EUR 85 million (Federal Republic of Germany 1989, § 10). The new Directive rules out financial ceilings in national law (Recital 56), and Germany's transposition act of 22 October 2026 removes the cap for products placed on the market from 9 December 2026 (Federal Republic of Germany 2026; Deutscher Bundestag 2026). The development-risk defence survives (Art. 11(1)(e)), although Member States may keep or introduce derogations from it (Art. 18); Germany keeps one only for genetic engineering (Deutscher Bundestag 2026). Other recent changes frame the reform. The proposed AI Liability Directive, which would have eased fault-based claims, was withdrawn in 2025 (European Commission 2025), and the AI Act's obligations for high-risk systems have been postponed to December 2027 and August 2028 (European Parliament and Council 2026), so that strict liability for AI arrives a year before the Act's rules for high-risk

systems, which are meant to prevent some of the harms it compensates.

Losses that the Directive does not cover, pure economic loss and damage to professional property and data, are not channelled upstream by it at all. Where the claim does return downstream and the insurer of the downstream manufacturer has adopted the generative-AI exclusion for products and completed operations (form CG 35 08, Section 2.4), the share that returns may itself be uninsured. Because all three cases in Section 2.3 fall outside the Directive, the harms it covers may be a minority of AI-related losses.

A.2 Prudential rules

In the Solvency II standard formula, premium and reserve risk for general liability is three times a volume measure scaled by standard deviations of 14% for premiums and 11% for reserves, aggregated with other lines through fixed correlations of 0.25 to 0.5 (European Commission 2015, Art. 115 and Annexes II and IV). Man-made catastrophe risk in liability is a scenario loss equal to a multiple of the premium expected over the following year: 100% for professional malpractice and for other general liability, 160% for employers' and for directors' and officers' liability, and 210% for non-proportional reinsurance (Art. 133 and Annex XI). None of these parameters depends on the technology that insureds share.

The nineteen critical providers designated under the Digital Operational Resilience Act in November 2025 include the cloud businesses of Amazon, Google and Microsoft (European Supervisory Authorities 2025; Rossa and Read 2025), groups that also develop or distribute foundation models. The designation covers the services they supply to financial institutions and leaves aside the dependence of those institutions' policyholders and borrowers on the same groups.

A.3 Who bore the CrowdStrike outage

Only 10 to 20% of the Fortune 500 loss was insured, because retentions were large and limits low relative to outage losses (Parametrix 2024). The vendor's standard terms cap its liability at the fees for the relevant offering and exclude lost profits (CrowdStrike 2026); Delta Air Lines, whose negotiated agreement capped liability at twice the fees except for gross negligence or wilful misconduct, has had to litigate through that exception, and two years after the outage its claim of at least USD 500 million was still in discovery (Delta Air Lines 2024; CrowdStrike Holdings 2026). The vendor's concessions took the form of subscription extensions, discounts and additional modules worth about USD 80 million of annual recurring revenue, about 1.5% of the Fortune 500 loss alone (CrowdStrike Holdings 2025). The largest cost outside the affected firms fell on the vendor's shareholders, who lost USD 25 billion of market value within eleven days (Stempel 2026). The vendor's share is a lower bound while Delta's claim is pending; if that claim succeeded in full, the share borne by the vendor and its insurers would rise to about a tenth of the Fortune 500 loss.

A.4 Providers' contract terms

Every provider whose terms could be read in full caps its liability to a business customer at about a year of fees, or, for an openly licensed model, excludes liability for consequential loss altogether, and all but one disclaim any warranty that outputs are accurate. The indemnities

Table 8. Who bore the loss from the CrowdStrike outage of 19 July 2024

Holder	What it bore	Source
Affected firms	Direct loss of the affected US Fortune 500 companies, excluding Microsoft, USD 5.4 billion; 80 to 90% uninsured	Parametrix (2024)
Insurers	10 to 20% of the Fortune 500 loss, USD 0.54 to 1.08 billion; global insured loss USD 0.3 to 1.5 billion	Parametrix (2024); CyberCube (2024); Guy Carpenter (2024)
Vendor, by contract	Standard terms: liability capped at the fees paid or payable for the offering during its subscription term, lost profits and revenue excluded. Agreement with Delta Air Lines: twice the fees, except for gross negligence or wilful misconduct	CrowdStrike (2026); CrowdStrike v. Delta (2024)
Vendor, in practice	Subscription extensions, discounts and extra modules worth about USD 80 million of annual recurring revenue in fiscal 2025; cash settlement offers immaterial and covered by the vendor's insurers	CrowdStrike Holdings (2025, 2026)
Litigation	Delta claims at least USD 500 million; gross negligence and computer trespass claims survived dismissal in May 2025; in discovery in 2026. Passengers' class action against the vendor dismissed, affirmed on appeal in May 2026	Delta Air Lines (2024); CrowdStrike Holdings (2026); Del Rio v. CrowdStrike (2026)
Vendor's shareholders	Market value down USD 25 billion within eleven days	Stempel (2026)

Notes: The loss and insured-loss figures are estimates made by modelling firms within weeks of the event, not tallies of paid claims. The insured share of 10 to 20% for the Fortune 500 reflects retentions, waiting periods and limits that were low relative to outage losses.

Table 9. What foundation-model providers promise their business customers

Provider and terms	Cap on the provider's liability	Provider indemnifies the customer against	Customer indemnifies the provider against
OpenAI (Services Agreement, 2026)	Amount paid in the 12 months before the event; not for gross negligence, wilful misconduct or indemnities	Claims that the services or outputs infringe intellectual property; output indemnity uncapped, with six exclusions	Claims related to the customer's applications or to its inputs and outputs
Anthropic (Commercial Terms, 2025)	Fees paid in the previous 12 months; not for indemnities	Claims that paid use, including training data, or outputs infringe intellectual property	Claims related to the customer's inputs or to use that breaches the usage policy
Google Cloud (Terms of Service, 2026)	Fees paid in the 12 months before the event; not for fraud, indemnities or infringement	Intellectual property claims on the services, on unmodified output of listed paid services and on training data	Claims arising from the customer's applications and data, which include generated output
Microsoft (Product Terms, 2026)	Amount paid in the 12 months before the claim under the legacy online subscription agreement	Copyright claims on output of covered products, if required mitigations are in place	Claims arising from "high-risk use", including claims based in strict liability or Microsoft's negligence
Amazon Web Services (Customer Agreement, 2026)	Amounts paid in the 12 months before the liability arose; in the EU, not for gross negligence or wilful misconduct	Intellectual property claims on the services; uncapped for output of Amazon's own models	Any third-party claim concerning the customer's or its end users' use of the services
Mistral AI (Commercial Terms, 2026)	Amounts paid in the 12 months before the event	Intellectual property claims on the products; no indemnity specific to outputs	Claims arising from the customer's own offerings that use the products
Cohere (SaaS Agreement)	Fees paid in the 12 months before the event; damages for personal injury or death and for inaccurate information excluded	Intellectual property claims on the services; copyright claims on outputs	Claims arising from customer data or breach
Meta (Llama 4 licence, 2025)	No liability for indirect or consequential damages under any theory, product liability included	None	Any third-party claim arising from use or distribution of the models

Notes: Published terms for business use of each provider's models, retrieved on 28 October 2026 (OpenAI 2026b,a; Anthropic 2025b; Google Cloud 2026; Microsoft 2026, 2019; Amazon Web Services 2026b,a; Mistral AI 2026; Cohere 2026; Meta 2025). Negotiated enterprise agreements may differ. Seven of the eight disclaim any warranty that outputs are accurate or error-free; Microsoft's generative-AI terms instead make the customer responsible for the applications it builds. For Microsoft the cap of the current customer agreement could not be retrieved; the legacy agreement is shown. For Amazon the output indemnity is described in the provider's documentation; the service terms that contain it could not be retrieved in full.

that providers give concern intellectual property: claims that outputs, or the data used to train the models, infringe the rights of others, often without a cap. None covers a claim that an output harmed a third party in any other way, through bodily injury, discrimination or misleading advice. In most of the terms the customer instead indemnifies the provider against third-party claims that arise from its applications, its inputs or the outputs it uses; Anthropic's customer indemnity is narrower, covering inputs and use in breach of its usage policy. Microsoft's terms place on the customer the risk of "high-risk use", any use in which a failure of the service could cause death, serious bodily injury or physical damage, and Cohere's exclude damages for personal injury or death. The terms of OpenAI and Anthropic also require individual arbitration and waive class proceedings, so that many customers harmed by the same failure would have to claim one by one. Between businesses, then, the share of a provider-level loss outside intellectual property that a deployer can pass to its provider by contract is bounded by a year of fees, which is small relative to the liability that a defective model could create for a large deployer (Appendix G.2).

B Further results and proofs

This appendix states the three results that Section 3 summarises, Lemma 2 on expected shortfall, Lemma 3 on value-at-risk and Proposition 3 on the scale of the accumulation, and proves the results of that section. Throughout, $F = (F_1, \dots, F_K)$ and $T = \sum_k \sigma_k F_k$. Two facts about majorisation are used repeatedly (Marshall et al. 2011, Ch. 3). A vector σ' majorises σ if, after ordering both in decreasing order, all partial sums of σ' are at least as large as those of σ, with equal totals. A function that is symmetric and convex on the simplex is Schur-convex, that is, non-decreasing in the majorisation order (Proposition 3.C.2 in Marshall et al. 2011).

Proof of Lemma 1. (i) F takes finitely many values. Conditional on F, the policies written on clients of provider k are independent and identically distributed with mean $c\mu(\pi + qF_k)$ and finite variance, so, for a nested sequence of books in which each book adds policies to the previous one, their average converges almost surely to that mean as $n_k = \sigma_k n \to \infty$ by the strong law of large numbers. Hence $\bar{Y}_n = \sum_k \sigma_k \big(n_k^{-1} \sum_{i \in k} Y_i\big) \to c\mu(\pi + qT)$ almost surely conditional on each value of F, and therefore unconditionally.

(ii) By the law of total variance, $\operatorname{Var}(\bar{Y}_n) = \operatorname{Var}\big(\mathrm{E}[\bar{Y}_n \mid F]\big) + \mathrm{E}\big[\operatorname{Var}(\bar{Y}_n \mid F)\big]$. The first term is $c^2\mu^2 q^2 \operatorname{Var}(T) = c^2\mu^2 q^2 \sum_k \sigma_k^2 p_k(1 - p_k)$, using the independence of the F_k. Conditional on F the policies are independent, so $\operatorname{Var}(\bar{Y}_n \mid F) = n^{-2} \sum_i \operatorname{Var}(Y_i \mid F)$, whose expectation is $c^2 V / n$ with $V = \sum_k \sigma_k \mathrm{E}[\operatorname{Var}(Y_i / c \mid F_k)]$ for i a client of k.

(iii) With $p_k = p$, $\operatorname{Var}(T) = p(1 - p)H(\sigma)$, which is also the variance of the average of $1/H$ independent Bernoulli(p) variables when $1/H$ is an integer. For a single policy, $Y_1 / c = XI + X'J$ with $I \sim \text{Bernoulli}(\pi)$, $J \sim \text{Bernoulli}(qp)$, I independent of J, and X, X' independent of both. Hence $\operatorname{Var}(Y_1/c) = \mu^2(1 + \delta^2)\pi - \mu^2\pi^2 + \mu^2(1 + \delta^2)qp - \mu^2 q^2 p^2$ and the covariance term vanishes. Since $\operatorname{Var}(\bar{Y}_n)$ decreases in n towards $c^2\mu^2 q^2 p(1 - p)H$, $n_{\text{eff}}(n)$ increases towards the ratio in (3). □

Lemma 2 (Expected shortfall and market structure)**.** *(i)* Monotonicity. *If $p_k = p$ for all k, then $\mathrm{ES}_\alpha(T)$ and $\operatorname{Var}(T)$ are Schur-convex in σ. If σ' majorises σ, that is, if market share has moved from smaller to larger providers, $\mathrm{ES}_\alpha(T(\sigma')) \geq \mathrm{ES}_\alpha(T(\sigma))$. Markets with different*

numbers of providers are compared by adding providers with zero share and the same failure probability.

(ii) Bounds. *For any* $p_1, \dots, p_K$,

$$\max_k \sigma_k \min\left\{1, \frac{p_k}{1-\alpha}\right\} \;\le\; \mathrm{ES}_\alpha(T) \;\le\; \min\left\{1, \frac{\bar{p}}{1-\alpha}\right\}. \tag{12}$$

Expected shortfall gets the ranking right because it is coherent (Artzner et al. 1999; Acerbi and Tasche 2002): it is convex and, with independent failures of equal probability, symmetric in the vector of shares, hence Schur-convex. It need not rise with the Herfindahl index when neither vector of shares majorises the other, and the bounds in (12) show why: capital is driven by the share of the largest provider, scaled by its failure probability relative to the tail probability.

Proof of Lemma 2. (i) Fix p and regard $T(\sigma) = \sum_k \sigma_k F_k$ as a function of σ on the simplex. Expected shortfall is positively homogeneous and subadditive (Acerbi and Tasche 2002). For σ, σ' and $\omega \in [0,1]$, $T(\omega\sigma + (1-\omega)\sigma') = \omega T(\sigma) + (1-\omega)T(\sigma')$, so $\mathrm{ES}_\alpha\big(T(\omega\sigma + (1-\omega)\sigma')\big) \le \omega\,\mathrm{ES}_\alpha(T(\sigma)) + (1-\omega)\,\mathrm{ES}_\alpha(T(\sigma'))$: the map $\sigma \mapsto \mathrm{ES}_\alpha(T(\sigma))$ is convex. When the F_k are independent and identically distributed, permuting σ does not change the distribution of $T(\sigma)$, and expected shortfall depends only on the distribution, so the map is symmetric. A symmetric convex function is Schur-convex. For the variance, $\mathrm{Var}(T) = p(1-p)\sum_k \sigma_k^2$, which is symmetric and convex.

(ii) For the lower bound, fix k. Because $T \ge \sigma_k F_k$, $\Pr(T < \sigma_k) \le 1 - p_k$, so $\mathrm{VaR}_u(T) \ge \sigma_k$ for every $u > 1 - p_k$. If $p_k \ge 1 - \alpha$, every $u \in (\alpha, 1)$ satisfies $u > 1 - p_k$, and $\mathrm{ES}_\alpha(T) \ge \sigma_k$. If $p_k < 1 - \alpha$, then, since $\mathrm{VaR}_u(T) \ge 0$,

$$\mathrm{ES}_\alpha(T) \ge \frac{1}{1-\alpha}\int_{1-p_k}^{1} \sigma_k \, du = \frac{\sigma_k p_k}{1-\alpha}.$$

Taking the maximum over k gives the lower bound. For the upper bound, $T \le 1$ gives $\mathrm{ES}_\alpha(T) \le 1$, and $\mathrm{VaR}_u(T) \ge 0$ gives $\int_\alpha^1 \mathrm{VaR}_u(T)\,du \le \int_0^1 \mathrm{VaR}_u(T)\,du = \mathrm{E}[T] = \bar{p}$. □

Lemma 3 (Value-at-risk and market structure)**.** *(i)* Tail blindness. *If the probability that at least one provider fails,* $1 - \prod_k (1 - p_k)$*, does not exceed* $1 - \alpha$*, then* $\mathrm{VaR}_\alpha(T) = 0$*, and capital under value-at-risk in* (4) *is zero for any shares.*

(ii) Non-monotonicity. *Value-at-risk is not monotone in concentration. If* $p_k = p$ *with* $p \le 1 - \alpha < 1 - (1-p)^K$*, a single provider serving all adopters yields* $\mathrm{VaR}_\alpha(T) = 0$*, while K equal providers yield* $\mathrm{VaR}_\alpha(T) \ge 1/K > 0$.

Proof. (i) $\Pr(T = 0) \ge \Pr(F = 0) = \prod_k (1 - p_k) \ge \alpha$, so $\inf\{t : \Pr(T \le t) \ge \alpha\} = 0$, and $[0 - \mathrm{E}[T]]^+ = 0$.

(ii) With a single provider, $\Pr(T = 0) = 1 - p \ge \alpha$, so $\mathrm{VaR}_\alpha(T) = 0$ by (i). With K equal providers, $\Pr(T = 0) = (1-p)^K < \alpha$, so the α-quantile is at least the smallest positive value of T, which is $1/K$. □

Applied to market structure, Lemma 3 implies that a stand-alone 1-in-200 calibration would set capital at zero for a market whose insureds all rely on one provider with a failure

probability of 0.4%, and at the failure of one provider for the same market served by ten. With seven providers, as in the baseline, part (i) applies only for failure probabilities below 0.07%, and for probabilities between 0.07% and 0.5% value-at-risk at 99.5% ranks a single provider as safer than seven equal ones. Appendix C shows that the reversal is confined to the stand-alone measure.

Proof of Proposition 1. (i) follows from the rearrangement in the text. (ii) With $p_k = p$, $\bar{p} = p$ does not depend on σ, so Ξ and $[\rho(T) - \bar{p}]^+$ are non-decreasing transformations of $\mathrm{ES}_\alpha(T)$, which is Schur-convex by Lemma 2(i); $\max_k \sigma_k$ is symmetric and convex, hence Schur-convex. The left-hand side of (6) is $a_w z$ times one of these two terms.

(iii) A planner who counts expenses, the cost of capital, the frictional costs v that cover avoids and the external cost Γ compares, per adopter, $\eta\mu h + \kappa\mu q[\rho(T) - \bar{p}]^+$ under cover with $v\mu h + \mathrm{E}[\Gamma(D_0(T))]/a$ under exclusion; the expected losses μh are borne in both cases. Dividing by $\kappa\mu q(r + \bar{p})$ gives the stated condition. For the last claim, $D_0(T)$ in (10) is affine and non-decreasing in T and Γ is convex and non-decreasing, so $\Gamma_D(t) = \Gamma(D_0(t))$ is convex. For every realisation of F the map $\sigma \mapsto \Gamma_D(\sum_k \sigma_k F_k)$ is convex, and expectations preserve convexity, so $\sigma \mapsto \mathrm{E}[\Gamma_D(\sum_k \sigma_k F_k)]$ is convex; it is symmetric when the F_k are independent and identically distributed. It is therefore Schur-convex, and so is W, whose denominator does not depend on σ when $p_k = p$. □

Proposition 3 (Scale of the accumulation)**.** *For the capitals e and $\bar{e}$ in* (8)*:*

(i) $\bar{e}(\theta)$ and $e(\theta)$ are non-decreasing in θ, and $0 \le \bar{e}(\theta) \le e(\theta) \le \rho(T) - \mathrm{E}[T]$.

(ii) As $\theta \to 0$, $\bar{e}(\theta)$ and $e(\theta)$ tend to zero; as $\theta \to \infty$, both tend to the stand-alone capital $\rho(T) - \mathrm{E}[T]$.

(iii) If $p_k = p$ for all k, $\bar{e}(\theta)$ is Schur-convex in σ for every θ. If $v > \eta$, the largest amount that the carrier covers at the private threshold,

$$\theta^* = \sup\{\theta : \bar{e}(\theta) \le (r + \bar{p})(v - \eta)/\kappa\},$$

is then positive, infinite if the stand-alone index Ξ is at most $(v - \eta)/\kappa$, and does not rise when σ becomes more concentrated in the majorisation order.

Proof. Write $f(\theta) = \rho(\theta T + R)$ for $\theta \ge 0$. Because expected shortfall is subadditive and positively homogeneous, $f(\omega\theta_1 + (1-\omega)\theta_2) \le \omega f(\theta_1) + (1-\omega) f(\theta_2)$, so f is convex. Since T takes finitely many values and is independent of R, which has a density, $\theta T + R$ has a density for every $\theta \ge 0$, and f is differentiable with $f'(\theta) = \mathrm{E}[T \mid \theta T + R \ge \mathrm{VaR}_\alpha(\theta T + R)]$ (Scaillet 2004). Hence $e(\theta) = f'(\theta) - \mathrm{E}[T]$, and because a differentiable convex function has a continuous derivative, $f(\theta) - f(0) = \int_0^\theta f'(u)\,du$, which gives $\bar{e}(\theta) = \theta^{-1} \int_0^\theta e(u)\,du$.

(i) The derivative of a convex function is non-decreasing, so e is non-decreasing; $\bar{e}$, the average of e over $[0, \theta]$, is then non-decreasing and does not exceed $e(\theta)$. For the lower bound, the event $\Omega = \{R \ge \mathrm{VaR}_\alpha(R)\}$ has probability $1 - \alpha$, and expected shortfall is at least the expectation conditional on any event of that probability, so $f(\theta) \ge \mathrm{E}[\theta T + R \mid \Omega] = \theta\,\mathrm{E}[T] + f(0)$ by the independence of T and R; hence $\bar{e} \ge 0$. For the upper bound, $e(\theta) + \mathrm{E}[T]$ is the expectation of T conditional on an event of probability $1 - \alpha$, which cannot exceed $\rho(T)$.

(ii) As $\theta \to 0$, the continuous derivative $f'(\theta)$ converges to $f'(0)$. Because R has a continuous distribution and is independent of T, $f'(0)$ is the expectation of T conditional on

the tail event of R alone, $\mathrm{E}[T \mid \Omega] = \mathrm{E}[T]$; hence $e(\theta) \to 0$, and $0 \le \bar{e} \le e$ gives the same limit for $\bar{e}$. As $\theta \to \infty$, positive homogeneity gives $(f(\theta) - f(0))/\theta = \rho(T + R/\theta) - \rho(R)/\theta$, which tends to $\rho(T)$ because $R/\theta \to 0$ in L^1 and expected shortfall is Lipschitz in L^1 with constant $1/(1-\alpha)$. So $\bar{e}(\theta) \to \rho(T) - \mathrm{E}[T]$, and $\bar{e} \le e \le \rho(T) - \mathrm{E}[T]$ gives the same limit for e.

(iii) For fixed θ, the map $\sigma \mapsto \rho(\theta \sum_k \sigma_k F_k + R)$ is convex, by the argument in the proof of Lemma 2(i), and symmetric when the F_k are independent and identically distributed and independent of R; it is therefore Schur-convex, and $\rho(R)$ and $\mathrm{E}[T] = p$ do not depend on σ. Because $\bar{e}$ is non-decreasing in θ, the set of amounts the carrier covers is an interval with upper end θ^*. It is non-degenerate because $\bar{e}(\theta) \to 0$ while the right-hand side of the condition is positive when $v > \eta$; it is unbounded if $\rho(T) - \mathrm{E}[T] \le (r + \bar{p})(v - \eta)/\kappa$, because $\bar{e}$ never exceeds that capital; and it does not grow when σ is replaced by a vector that majorises it, because $\bar{e}$ does not fall. □

Proof of Proposition 2. Write $f_1(u) = \rho(uT + U)$ for $u \ge 0$, with U distributed as the U_j. By positive homogeneity, $\rho(\theta T + s_j U_j) = s_j f_1(\theta/s_j)$, so carrier j's capital for θ is $\kappa[s_j f_1(\theta/s_j) - s_j f_1(0) - \theta\,\mathrm{E}[T]]$. By the proof of Proposition 3, f_1 is convex and differentiable with $f_1'(u) - \mathrm{E}[T] = e_1(u)$; a differentiable convex function has a continuous derivative, and $e_1(u) \to 0$ as $u \to 0$, so e_1 is continuous and non-decreasing on $[0, \infty)$, with $e_1(0) = 0$ and limit $\rho(T) - \mathrm{E}[T]$.

(i) Carrier j maximises $\Pi_j(\theta) = m\theta - \kappa[s_j f_1(\theta/s_j) - s_j f_1(0) - \theta\,\mathrm{E}[T]]$, which is concave with derivative $m - \kappa e_1(\theta/s_j)$. Any θ with $e_1(\theta/s_j) = m/\kappa$ is optimal, and $\theta_j = s_j g(m/\kappa)$ is the largest such choice; for $m \ge \kappa[\rho(T) - \mathrm{E}[T]]$, g is infinite.

(ii) At a margin $0 < m < \kappa[\rho(T) - \mathrm{E}[T]]$, carrier j's optimal holdings form the interval $s_j[g_-(m/\kappa), g(m/\kappa)]$, where $g_-(t) = \inf\{u \ge 0 : e_1(u) \ge t\}$, because e_1 is continuous and non-decreasing. Aggregate holdings can equal one if and only if $1/S \in [g_-(m/\kappa), g(m/\kappa)]$, that is, if and only if $e_1(1/S) = m/\kappa$. The clearing margin is therefore unique, $m^* = \kappa e_1(1/S)$, and the allocation $\theta_j = s_j/S$ clears the market at it. Insureds buy full cover if and only if the margin does not exceed $\bar{m}$, the condition (5) with the marginal cost of the accumulation in place of its stand-alone cost.

(iii) Let $y_u = \mathrm{VaR}_\alpha(uT + U)$ and $\bar{\Phi} = 1 - \Phi$, the standard normal survival function. Since $\Pr(uT + U \ge y_u) = 1 - \alpha$,

$$e_1(u) = \frac{1}{1-\alpha} \sum_t \Pr(T = t)\,(t - \mathrm{E}[T])\,\bar{\Phi}(y_u - ut).$$

A first-order expansion, $\bar{\Phi}(y_u - ut) = \bar{\Phi}(y_u) + ut\,\phi(y_u) + O(u^2)$, uniformly over the finitely many values t, and $\sum_t \Pr(T = t)(t - \mathrm{E}[T]) = 0$ give $e_1(u) = u\,\mathrm{Var}(T)\phi(y_u)/(1-\alpha) + O(u^2)$. As $u \to 0$, $y_u \to \Phi^{-1}(\alpha)$, so $e_1(u) = u\,\mathrm{Var}(T)\phi(\Phi^{-1}(\alpha))/(1-\alpha) + o(u)$. With $p_k = p$, $\mathrm{Var}(T) = p(1-p)H(\sigma)$, and $u = 1/S$ gives (9). As $S \to 0$, $1/S \to \infty$ and $e_1(1/S) \to \rho(T) - \mathrm{E}[T] = \rho(T) - \bar{p}$ by Proposition 3(ii).

(iv) Total holdings cannot exceed $\sum_j \bar{\theta}_j$.

(v) Carrier j covers its own accumulation if and only if its incremental capital per unit satisfies $\kappa \bar{e}_j(w_j) \le \bar{m}$, and $\bar{e}_j(w_j) = \bar{e}_1(w_j/s_j)$ by positive homogeneity. Because $\bar{e}_1$ is non-decreasing, this holds if and only if $w_j/s_j \le \theta_1^*$. By Proposition 3(iii), θ_1^* does not rise when σ is replaced by a vector that majorises it, so the set of carriers with $w_j/s_j > \theta_1^*$ does not shrink; scaling every w_j/s_j up has the same effect. The excluded share is a sum of non-negative terms over that set. □

Proof of Corollary 1. (i) follows from Proposition 2(v).

(ii) $\Psi(c) = \Pr((1-c)X > B)$ is non-increasing in c, so withdrawing cover raises it. H and $\max_k \sigma_k$ are symmetric and convex, hence Schur-convex, so $H(\sigma') \geq H(\sigma)$ and $\max_k \sigma'_k \geq \max_k \sigma_k$. Both quantities are products of non-negative factors that do not fall from σ to σ'. The expression for the variance follows from (10) and $\mathrm{Var}(T) = p(1-p)H$. □

C Calibration details

This appendix sets out the parameters and results summarised in Section 4 (Sections C.1 to C.3) and the calculations behind them (Section C.4).

C.1 Parameters

Provider shares are Menlo Ventures' end-2025 shares, 40, 27 and 21% for the three largest providers, with the residual split equally among four providers (Tully et al. 2025b); the other sources imply more concentration (Section 2.1), and shares of spending are only proxies for the share of production workflows that depend on each provider.

The event record gives two output defects that reached interface customers in 13.0 to 15.3 provider-years, depending on the date from which Meta is counted (Section 2.2). If between one in ten and one in a hundred such defects produced harm at scale, the annual probability p of a liability-relevant failure would lie between 0.13 and 1.5%, with a geometric midpoint of 0.4 to 0.5%, and the exact Poisson interval of the rate widens the range to 0.016 to 5.6% (Appendix D). I take $p = 0.4\%$ as the base case and $p = 1\%$ as a stress case, the same for all providers, and figures shade the range from 0.13 to 1.5%. The stress value equals $1 - \alpha$ for expected shortfall at 99% and maximises the stand-alone index of the end-2025 market, so it is the least favourable value for cover assessed stand-alone; the shared price, roughly proportional to $p(1-p)$ in (9), keeps rising above it.

The premium-weighted adoption rate a_w, the share of liability premium written on firms whose operations rely on a foundation model, is 0.20, 0.35 or 0.55. In 2025, 32% of EU enterprises with 250 or more persons employed and 9% of all enterprises used AI that generates language (Eurostat 2026a); because large firms buy most liability cover, the premium-weighted rate lies close to the large-firm share, and the central value allows for a year of further growth. The values of 0.20 and 0.55 are illustrative bounds; for comparison, 20% of all enterprises and 55% of large ones used some form of AI in 2025 (Eurostat 2025). Corporate credit is concentrated on larger firms in the same way, so 35% is also the central share of corporate loans that go to adopters, probably a high value, since the use of AI that generates language includes uses that create no liability and corporate credit includes loans to small firms, where adoption is lower. Bank losses are proportional to this share and to the share of borrowers whose AI-related liability is excluded, which the bank figures set at one. They are therefore losses per unit of that share, an overstatement for Europe today, since every documented exclusion is in a US or Lloyd's wording. The insurance aggregates are American and the adoption and banking data European; because the results depend on ratios, the mix has little effect.

The event intensity $z = q\mu / P_i$, the expected AI-related liability loss of a client of a failing provider in units of its annual liability premium P_i, runs from 0.5 to 4; at $z = 1$ and $q = 10\%$

a harm costs ten years' premium. Because z is scaled on US liability premium, it reflects American claims practice, with class actions and punitive damages. The central scenario is severe: with US liability cover silent on AI, it would cost about USD 20 billion of insured losses (Appendix C.2), more than ten times the highest estimate for the CrowdStrike outage, USD 1.5 billion, all of it from downtime (Table 8 in Appendix A). Bank losses scale with $q\Psi_0$, the probability that a client of the failing provider is harmed and defaults. The default probability Ψ_0 of a harmed firm without cover is 15 or 30% at $z = 1$ and $q = 10\%$; tied to the intensity through a lognormal loss per harm and a default buffer of 16.8 years of liability premium, chosen rather than measured, it is 3, 15, 42 and 74% at z of 0.5, 1, 2 and 4 (Appendix C.4), so if European claims per harm were half the US-scaled intensity, the central bank figures would fall to a fifth. Large firms, whose equity is many times their liability premium, would default less often than this buffer implies. The exclusion thresholds $(v - \eta)/\kappa$ are 2.5 and 5, a surplus from insurance, net of expenses, of 15 and 30% of expected losses at a 6% cost of capital, with $r = \pi/q = 0.1$, and a carrier's other risks are an independent normal loss with a standard deviation of 15 or 30% of premium.

C.2 Diversification, event losses and capital

At the end of 2025 the Herfindahl index of provider shares was 0.28, or 2,800 on the scale of competition policy, well above the 1,800 at which the 2023 US Merger Guidelines regard a market as highly concentrated (U.S. Department of Justice and Federal Trade Commission 2023), and a liability book that mirrors the market carries the AI risk of at most about 1,850 independent policies at $p = 0.4\%$ and 790 at $p = 1\%$, however large it is (Appendix C.4). That ceiling is held up by the idiosyncratic harms that accompany the common ones; for the provider factor alone it is 72, which is the number that bears on accumulation. A failure of the largest provider costs $a_w \sigma_{\max} z$ of a year's liability premium, 14% in the central scenario and 4 to 88% across the grid (Table 15), against a Solvency II catastrophe charge for general liability of 100% of a year's premium meant to cover every man-made liability catastrophe. On the USD 139.5 billion of US premiums in other liability and products liability (National Association of Insurance Commissioners 2026b), cover silent on AI would lose about USD 20 billion before reinsurance, 1.5% of the industry's surplus (National Association of Insurance Commissioners 2026a), and USD 6 to 123 billion across the grid, a large loss for the liability lines but not a threat to the sector's solvency. Stand-alone capital under expected shortfall is close to the share of the largest provider, about $\sigma_{\max}(1 - p)$ when $p = 1 - \alpha$: 0.40 of a full-market event at $p = 1\%$ and 0.31 at $p = 0.4\%$. Value-at-risk requires no capital at all for a market served by one provider at p of 0.25 or 0.4% (Appendix C.4).

C.3 Who can hold the accumulation

At the base and stress values, stand-alone capital would justify exclusion at the lower threshold but not at the upper one. Under expected shortfall at 99%, the stand-alone index Ξ for end-2025 shares is approximately $\sigma_{\max}(1 - p)/(r + p)$ near $p = 1\%$, where it peaks at 3.6; it is 3.0 at $p = 0.4\%$, crosses the lower threshold at $p \approx 0.29\%$ and is 1.3 at the lower end of the range implied by the event record (Figure 7). Cover survives only if insureds' surplus from it, net of expenses, is at least 18% of expected losses in the base case and 22% in the stress case. The

index rises to 6.6 with $r = 0.05$ and falls to 1.9 with $r = 0.2$; with the transaction-based shares it is 5.4 at $p = 1\%$, and the share of 56% reported for the largest provider in early 2026 bounds it from below by 5.0 (Lemma 2(ii)).

The same accumulation needs little capital inside a diversified book. For a carrier whose book mirrors the market and whose other risks have a standard deviation of 15% of premium, the accumulation is $a_w z/0.15 = 2.3$ times that standard deviation in the central scenario, and the index is 0.16 in the base case and 0.37 in the stress case under the Euler allocation, and 0.06 and 0.14 under incremental capital (Table 17 and Figure 6a). Shared across the US property-casualty sector, whose USD 1.1 trillion of direct premiums (National Association of Insurance Commissioners 2026a) put the summed standard deviation of other risks at $S = 3.4$ full-market events, the equilibrium index of Proposition 2 is 0.01 in the base case and 0.02 in the stress case, less than a hundredth of the lower threshold. At $p = 1\%$ its price rises roughly in proportion to H, from 0.01 with seven equal providers to 0.04 with the transaction-based shares and 0.10 with a single provider, values that the first-order formula (9) understates by about a sixth and three-tenths at this capacity. The single-provider value is also what a reinsurer that cannot observe which provider insureds rely on can charge to protect itself, since in numerical checks over random market structures no structure carries a higher price. Priced that way, the accumulation carries an index of 0.10 at the central exposure and 2.43 at the largest exposure on the grid ($a_w = 0.55$, $z = 4$), just below the lower threshold, so unobserved provider mixes do not by themselves keep the accumulation with the carriers that write it, except possibly at the most intense exposures.

If the sharing carriers' other risks load on the provider factor, with $R = s\big[\sqrt{1-\zeta^2}\,U + \zeta\,(T - \mathrm{E}[T])/\mathrm{sd}(T)\big]$, the sector-wide index at $p = 1\%$ is 0.13 at $\zeta = 0.05$ and 0.33 at $\zeta = 0.1$, and it reaches the lower threshold only at $\zeta = 0.34$, a loading at which the rest of the sector's book loses 2.5 standard deviations whenever the largest provider fails, as it might through AI-related assets. Aggregate appetites do not bind either: global reinsurance capital of USD 785 billion at the end of 2025 (Gallin 2026), with an appetite of a tenth of it for a single scenario, would absorb a largest-provider event of about USD 78 billion, four times the central scenario for silent US cover. If the accumulation stays with the carriers that write it, the reason lies in the frictions of Section 3.5.

In autarky, a carrier excludes at the lower threshold once the accumulation it holds exceeds 21 times the standard deviation of its other risks under incremental capital and 7 times under the Euler allocation (8 and 4 with the transaction-based shares; with seven equal providers, never). Across the 24 cells of the grid for a book that mirrors the market, the index under incremental capital never exceeds 2.1, and under the Euler allocation it exceeds 2.5 only in four cells, all at intensities of two or more (Table 17), so capital costs alone lead to exclusion only in books concentrated on AI-intensive insureds. Scenario appetites bind sooner, because diversification does not relax a limit on the loss of a scenario. With an appetite of a tenth of equity for the largest-provider scenario, a liability writer with equity of 60% of premium breaches it once $a_w z$ exceeds 0.15 (Figure 6b), and its scenario loss is 23% of equity in the central case. A carrier with the balance sheet of the US sector average, surplus of 115% of premium and 4.4% of premium written on adopters' liability, would lose 1.5% of its surplus at $z = 1$ and breach the appetite only at $z = 6.5$. The mechanism thus bears on specialty and liability writers concentrated on AI-intensive insureds, where affirmative cover is being written; the

multi-line groups that have filed most of the standard exclusions are too diversified for it to bind. The appetite and the equity ratio are illustrative values, which makes the exclusions they imply illustrative as well.

The largest provider's share and exclusions have moved in opposite directions: with 2023 shares, when the largest provider held half of the market, the stand-alone index at $p = 1\%$ was 4.5, while the share of large EU enterprises using AI that generates language quadrupled between 2023 and 2025, and the standard exclusions spread after this growth (Appendix C.4). If accumulation plays a part in the exclusions, it does so through adoption. Limits per provider event relax both constraints: a limit of a tenth of a full-market event per failing provider lowers the stand-alone index for the end-2025 market from 3.6 to 1.0 and the scenario loss to 5.8% of a liability writer's equity, while 95% of the expected AI loss remains insured and the insured firms retain three-quarters of the loss when the largest provider fails (Appendix F).

C.4 Supporting calculations

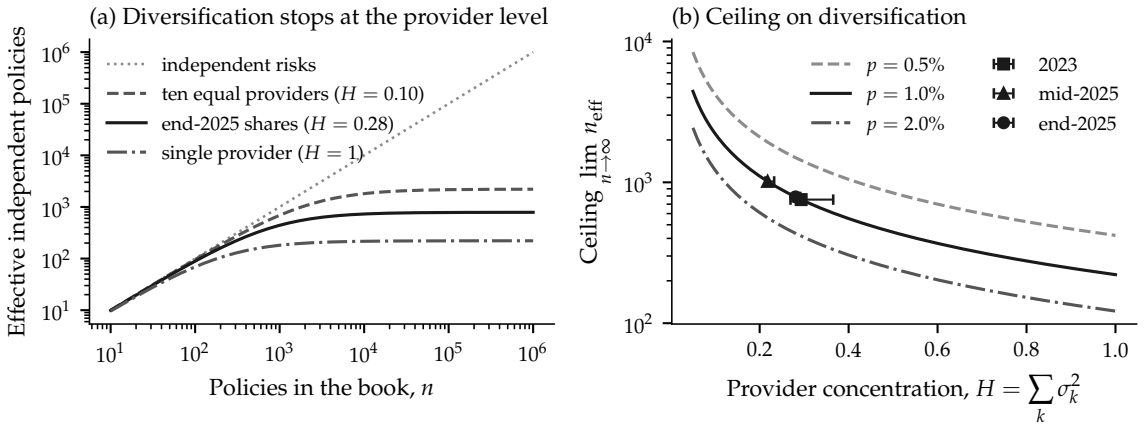


Figure 4. Diversification of the AI component of a liability book

Notes: Panel (a) plots $n_{\text{eff}}(n) = \text{Var}(Y_1)/\text{Var}(\bar{Y}_n)$ against the number of policies n for a book that mirrors the market. Panel (b) plots the ceiling (3) against H; markers show the three market snapshots with the range of H from Table 14. Parameters: $q = 10\%$, $\pi = 1\%$, coefficient of variation of severity one, $p = 1\%$ unless stated.

Who holds the provider factor. Figure 8 combines the two sides of Corollary 1 along a path of market structures that runs from seven equal providers through the end-2025 market to a single provider, for carriers that assess the accumulation stand-alone, the limit of autarky for books dominated by AI-exposed insureds, with $p = 1\%$, $q = 10\%$, $r = 0.1$ and, for corporate borrowers, $\Psi_0 = 15\%$. Such carriers withdraw cover once the largest provider holds 28% of the market if the threshold $(v - \eta)/\kappa$ is 2.5, and 56% if it is 5. Up to that point the systematic risk they hold grows with concentration; beyond it they hold none. For lenders the pattern is reversed: the failure-year loss of a lender whose borrowers mirror the market appears when cover is withdrawn and then grows with the share of the largest provider, from 17 basis points of exposure at the lower threshold to 24 at end-2025 shares and 60 for a single provider, whereas capital under the internal ratings-based formula rises once, by 34.5 basis points, mostly for idiosyncratic harms, and does not respond to concentration afterwards.

Table 10. Parameters of the quantitative illustration

Parameter	Symbol	Values	Source or rationale
Provider shares, end-2025	σ	0.40, 0.27, 0.21; residual 0.12	Tully et al. (2025b); residual split in four
Provider shares, mid-2025	σ	0.32, 0.25, 0.20, 0.09; residual 0.14	Tully et al. (2025a)
Provider shares, 2023	σ	0.50, 0.12, 0.07; residual 0.31	Tully et al. (2025b)
Provider shares, transaction data, week to 23 October 2026	σ	0.605, 0.351, 0.022, 0.019, 0.003	Ramp (2026); $H = 0.49$
Largest share, survey of large firms, early 2026	$\sigma_{\max}$	0.56	Wang et al. (2026); used only for the bound of Lemma 2(ii)
Provider failure probability	p	0.13–1.5%; base 0.4%, stress 1%; 0.25 and 2% in robustness checks	Event record and conversion range (Appendix C.1)
Harm probability in a failure	q	5, 10, 20%; central 10%	Grid
Idiosyncratic harm	$\pi = rq$	$r = 0.1$; 0.05 and 0.2 as alternatives	Grid
Coefficient of variation of severity	δ	1; 2 as robustness	Grid
Premium-weighted adoption	a_w	0.20, 0.35, 0.55	Central: language generation by 32% of large and 9% of all enterprises, premium-weighted, with growth; bounds illustrative, 55% of large enterprises used any AI (Eurostat 2025, 2026a)
Event intensity	z	0.5, 1, 2, 4	Grid; Appendix C.1
Risk measures	ρ	$\text{VaR}_{99.5\%}$, $\text{ES}_{99\%}$, $\text{ES}_{99.5\%}$	Solvency II; Swiss Solvency Test
Net surplus over cost of capital	$(v-\eta)/\kappa$	2.5, 5	$\kappa = 6\%$, $v - \eta = 15$ or 30% of expected loss
Rest of the insurer's book, s.d.		15, 30% of premium	Premium risk; premium and reserve risk
Scenario appetite, share of equity	ϵ	10%; equity of 60% of premium	Illustrative
Global reinsurance capital, end-2025		USD 785 billion; a tenth for one scenario	Gallin (2026); appetite illustrative
Probability of default	PD	1% corporate; 2% SME	Illustrative
Loss given default	ℓ	40%	Foundation IRB, senior unsecured
Asset correlation		Basel corporate function; SME adjustment at EUR 10m turnover	Basel Committee on Banking Supervision (2005)
Default probability given uninsured harm	Ψ_0	15, 30% at $z = 1$, $q = 10\%$	Grid; tied to z through the buffer B
Default buffer of a harmed firm	B	16.8 or 10.9 years of liability premium	Sets $\Psi_0 = 15$ or 30% at $z = 1$, $q = 10\%$, $\delta = 1$
AI infrastructure, share of a bank's corporate book	ι	1, 2, 5%	Grid; not published for any bank (Section 5.3)
Loss rate on AI-infrastructure exposure in the failure year	ϱ	12.5%; 5 to 25%	Illustrative; a default rate on facilities serviced by the failing provider times a loss given default on collateral whose value depends on the same demand
Share of claims allocated upstream	λ	0.2; 0.7 as an upper case	Contract terms; Appendix G
Claims expansion under the PLD	β	1.25, 1.5	Illustrative
Provider capacity, share of full-market claims	$\hat{A}$	2, 5, 10%	Illustrative

Table 11. Artificial intelligence in US annual reports

	Filings			Per cent of all filings		
	2024	2025	2026	2024	2025	2026
All 10-K filings	7,178	6,876	6,516	–	–	–
Mention artificial intelligence	2,436	3,324	3,901	33.9	48.3	59.9
Mention generative AI	728	1,223	1,599	10.1	17.8	24.5
Refer to third-party AI	79	214	454	1.1	3.1	7.0
of which name OpenAI, ChatGPT or Anthropic	–	–	41	–	–	0.6
Name a foundation-model provider	91	107	159	1.3	1.6	2.4

Notes: Form 10-K filings returned by EDGAR full text search for filing dates in each calendar year; 2026 runs from 1 January to 27 October (U.S. Securities and Exchange Commission 2026). "All 10-K filings" counts filings that contain the phrase "risk factors". Generative AI: "generative AI" or "generative artificial intelligence". Third-party AI: the phrase "third-party AI". Foundation-model provider: "OpenAI", "ChatGPT", "Anthropic", "Google Gemini" or "Cohere"; the count includes developers, investors, partners and competitors that name each other. Counts are of documents and include amended filings; of the first 100 hits for third-party AI in 2026, 99 were annual reports or amendments and all came from distinct filings. Of the 2026 filings referring to third-party AI that carry an industry code, at most 77% came from firms outside software and information technology services (SIC 7370 to 7374). The count of third-party AI filings that name OpenAI, ChatGPT or Anthropic is obtained by inclusion and exclusion from counts of joint mentions and is available for 2026 only. The counts measure disclosure, which can understate or overstate actual use.

Table 12. Administrative agents of credit facilities to firms that name a provider

Administrative agent	Facilities	Share (%)
JPMorgan Chase	11	35.5
Bank of America	5	16.1
Morgan Stanley	4	12.9
Wells Fargo	2	6.5
PNC	2	6.5
First Citizens (SVB)	2	6.5
Citigroup	1	3.2
Golub Capital†	1	3.2
Goldman Sachs	1	3.2
Ares Capital†	1	3.2
KeyBank	1	3.2
Total with an identified agent	31	100.0

Notes: Of the 93 firms whose filings name OpenAI or Anthropic as a dependency between January 2025 and October 2026, 32 disclose a credit facility and 31 name its administrative agent. One facility per firm, the most recent identified, so the counts are of firms and not of amounts; agents are grouped to the banking or credit group behind the contracting entity. †Private credit funds rather than banks. The administrative agent arranges and administers a syndicate and does not disclose the share of the loan it retains, and the other members are rarely named, so these counts measure arranging relationships and not held exposure. The firms are US filers and their agents are mostly US banks, a different population from the banks of Table 20.

Table 13. Administrative agents on the two channels

Administrative agent	Firms naming a provider	AI infra-structure	Total
JPMorgan Chase	11	6	17
Morgan Stanley	4	5	9
Bank of America	5	–	5
First Citizens (SVB)	2	–	2
Goldman Sachs	1	1	2
MUFG	–	2	2
PNC	2	–	2
Wells Fargo	2	–	2
Citigroup	1	–	1
First National Bank of Omaha	–	1	1
KeyBank	1	–	1
SMBC	–	1	1
Total	29	16	45

Notes: Counts of credit facilities by administrative agent. "Firms naming a provider": the 31 facilities of Table 12, of which 29 have a commercial bank as agent and two a private credit fund. "AI infrastructure": facilities of ten US-listed AI-infrastructure firms (CoreWeave, Applied Digital, Nebius, IREN, Cipher Mining, TeraWulf, Core Scientific, Galaxy Digital, Hut 8 and DigitalBridge), read from their filings with the US Securities and Exchange Commission; 31 financings are identified, 22 loan facilities of USD 33.0 billion and nine note issues of USD 16.9 billion. Only loan facilities with a commercial bank as administrative agent are counted, which leaves out fifteen of the 31 financings: ten administered by corporate trustees or loan-agency firms that take no credit risk, two extended by non-bank lenders, two note issues agented by a bank, and one term loan whose agent the filing does not name. Amounts are commitments at signing, not drawn balances. As in Table 12, an agent does not disclose the share it retains and the other syndicate members are not named, so the counts measure arranging relationships and not held exposure.

Table 14. Concentration of the provider market and the ceiling on diversification

Snapshot	Largest share	H (range)	$1/H$ (range)	n_{eff} ceiling
2023	0.50	0.27–0.37	2.7–3.7	754
mid-2025	0.32	0.21–0.23	4.3–4.7	1015
end-2025	0.40	0.28–0.29	3.4–3.6	788

Notes: Shares from Tully et al. (2025a,b); 2023 and end-2025 are shares of enterprise API spending, mid-2025 shares of enterprise usage. The range of H runs from a fully fragmented residual to a residual held by a single provider. The ceiling is $\bar{n}_{\text{eff}}$ in (3) for a book that mirrors the market and the baseline split of the residual, with $q = 10\%$, $p = 1\%$, $\pi = 1\%$ and a coefficient of variation of severity of one. With a coefficient of variation of two, the end-2025 ceiling is 1,980, and without idiosyncratic harms it is 72 (180 with a coefficient of variation of two).

Table 15. Loss from a failure of the largest provider, per cent of a year's liability premium

	Event intensity z			
Premium-weighted adoption	0.5	1	2	4
$a_w = 0.20$	4	8	16	32
$a_w = 0.35$	7	14	28	56
$a_w = 0.55$	11	22	44	88

Notes: Loss $= a_w \sigma_{\max} z$ with $\sigma_{\max} = 0.40$ (end-2025 share of the largest provider). a_w is the share of liability premium written on AI adopters; z is the expected AI-related liability loss of a client of the failing provider in units of its annual premium.

Table 16. Stand-alone capital for provider accumulation under three risk measures

		Annual failure probability p			
Market structure	Measure	0.25%	0.4%	1%	2%
Single provider	$\text{VaR}_{99.5\%}$	0.00	0.00	0.99	0.98
	$\text{ES}_{99\%}$	0.25	0.40	0.99	0.98
	$\text{ES}_{99.5\%}$	0.50	0.80	0.99	0.98
End-2025 shares	$\text{VaR}_{99.5\%}$	0.24	0.27	0.39	0.38
	$\text{ES}_{99\%}$	0.23	0.31	0.40	0.41
	$\text{ES}_{99.5\%}$	0.33	0.37	0.40	0.43
Ten equal providers	$\text{VaR}_{99.5\%}$	0.10	0.10	0.09	0.18
	$\text{ES}_{99\%}$	0.10	0.10	0.13	0.19
	$\text{ES}_{99.5\%}$	0.10	0.11	0.18	0.20

Notes: Capital $\rho(T) - \text{E}[T]$ in units of a full-market event μq per adopter, floored at zero for value-at-risk. End-2025 shares with the residual split among four providers. All providers have the same failure probability p.

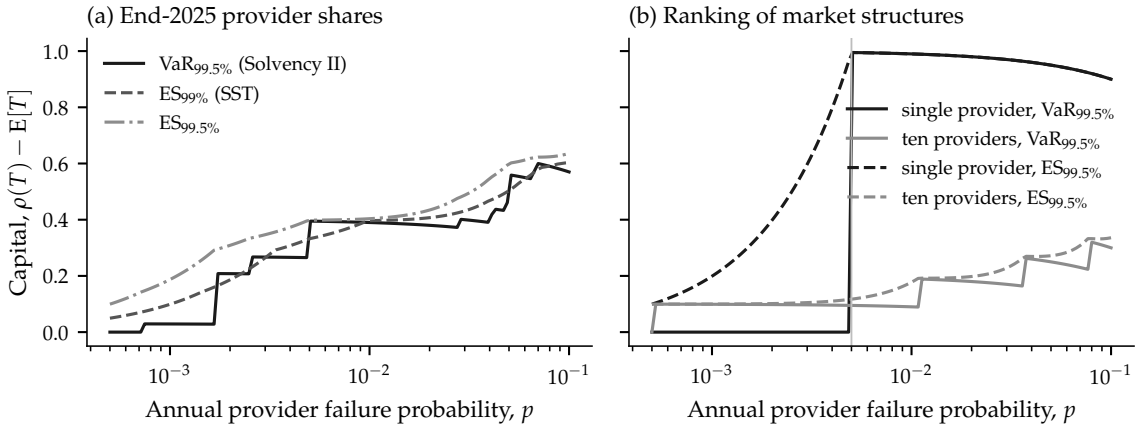


Figure 5. Stand-alone capital for provider accumulation as a function of the failure probability
Notes: Capital $\rho(T) - \text{E}[T]$ in units of a full-market event. Panel (a): end-2025 shares under three measures. Panel (b): a single provider and ten equal providers under value-at-risk and expected shortfall at 99.5%; the vertical line marks $p = 1 - \alpha = 0.5\%$.

Table 17. Exclusion index under integrated capital

	Rest of book: s.d. 15% of premium				Rest of book: s.d. 30% of premium			
Adoption	$z = 0.5$	$z = 1$	$z = 2$	$z = 4$	$z = 0.5$	$z = 1$	$z = 2$	$z = 4$
Incremental capital								
$a_w = 0.20$	0.03	0.06	0.18	0.59	0.01	0.03	0.06	0.18
$a_w = 0.35$	0.05	0.14	0.47	1.40	0.02	0.05	0.14	0.47
$a_w = 0.55$	0.10	0.30	1.00	2.08	0.04	0.10	0.30	1.00
Euler allocation								
$a_w = 0.20$	0.06	0.15	0.47	1.63	0.03	0.06	0.15	0.47
$a_w = 0.35$	0.12	0.37	1.30	3.07	0.05	0.12	0.37	1.30
$a_w = 0.55$	0.24	0.84	2.54	3.38	0.09	0.24	0.84	2.54
Stand-alone	3.61 in every cell							

Notes: Ξ for end-2025 shares, $p = 1\%$ (the stress case) and $r = 0.1$, when the AI accumulation is added to an independent normal loss whose standard deviation is 15 or 30% of a year's liability premium. Incremental capital is the increase in expected shortfall at 99% of the whole book; the Euler allocation is the accumulation's contribution to that expected shortfall, $\mathrm{E}[T \mid \text{whole-book loss in the 1\% tail}] - \mathrm{E}[T]$. Cover is supplied when $\Xi \leq (v - \eta)/\kappa$, between 2.5 and 5 at a 6% cost of capital. At $p = 0.4\%$, the base case, the central cell ($a_w = 0.35$, $z = 1$, standard deviation of 15% of premium) is 0.06 under incremental capital and 0.16 under the Euler allocation.

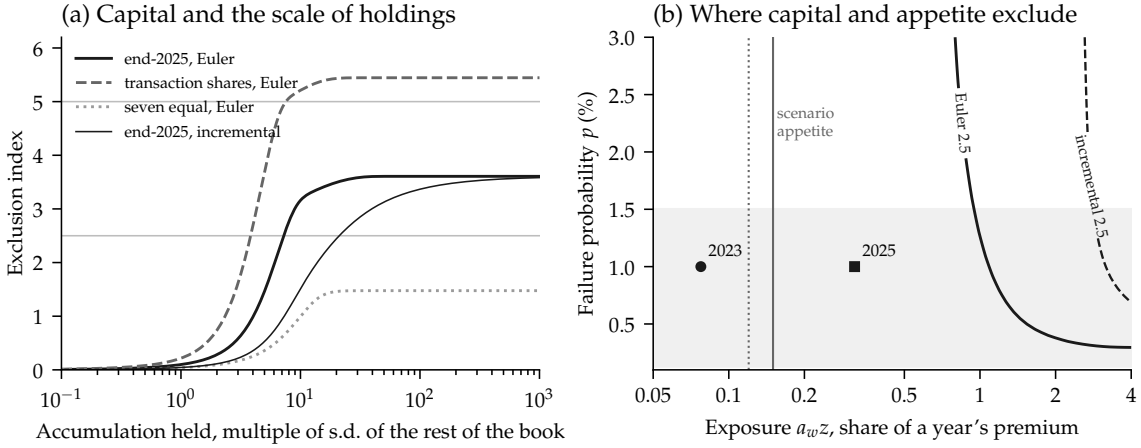


Figure 6. Capital for the accumulation and the scale at which it is held

Notes: Panel (a): exclusion index under the Euler allocation, $e/(r + \bar{p})$, and under incremental capital, $\bar{e}/(r + \bar{p})$, from (8), for a carrier that holds an accumulation of the size shown on the horizontal axis alongside an independent normal loss with unit standard deviation; $p = 1\%$, the stress case. The sharing equilibrium of Proposition 2 is the point $1/S$ on the horizontal axis. Horizontal lines mark the thresholds 2.5 and 5. Panel (b): contours at which the Euler index (solid) and the incremental index (dashed) of a book that mirrors the market reach the thresholds, as functions of the full-market exposure $a_w z$ and the failure probability, end-2025 shares, rest of the book with a standard deviation of 15% of premium; cover is excluded to the right of a contour. The vertical lines mark the scenario appetite, a loss from a failure of the largest provider of a tenth of equity with equity of 60% of premium, at end-2025 shares (solid) and 2023 shares (dotted); exposures to the right breach it. Markers show the exposure of 2023 and 2025 at $z = 1$, with a_w set to the share of large EU enterprises using AI that generates language. The shaded band is the range of p implied by the event record, 0.13 to 1.5%. Expected shortfall at 99% and $r = 0.1$ throughout.

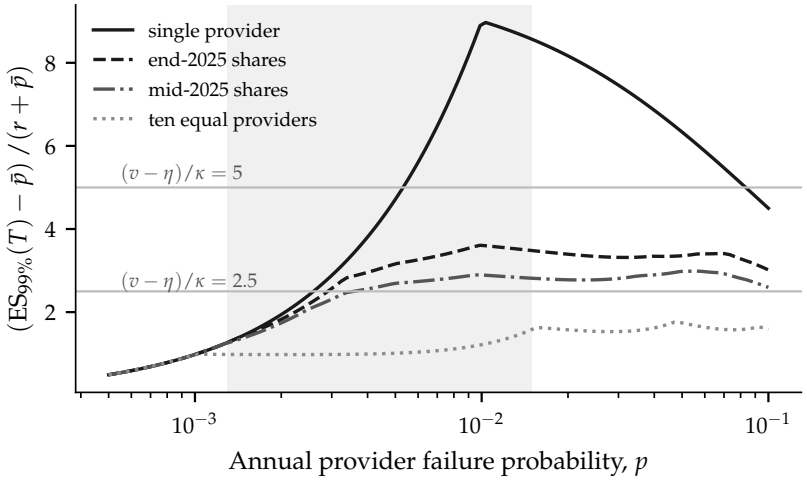


Figure 7. The exclusion frontier under stand-alone capital

Notes: $\Xi = (\mathrm{ES}_{99\%}(T) - \bar{p})/(r + \bar{p})$ with $r = \pi/q = 0.1$. Cover is supplied when $\Xi \leq (v - \eta)/\kappa$; horizontal lines mark the two values used in the text. The shaded band is the range of p implied by the event record, 0.13 to 1.5%.

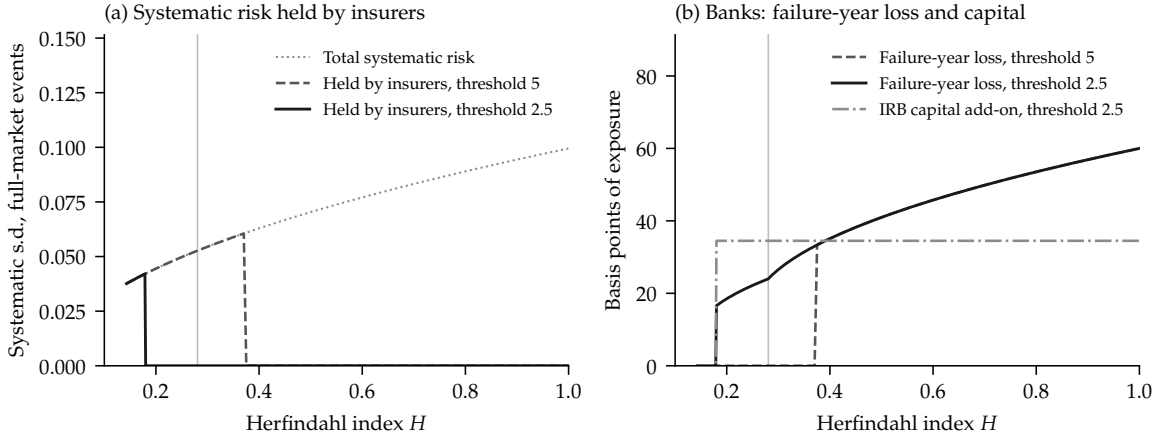


Figure 8. Who holds the provider factor as concentration rises

Notes: Market structures on a path from seven equal providers ($H = 0.14$) through the end-2025 market ($H = 0.28$, vertical line) to a single provider ($H = 1$), linear in each leg. Cover is supplied when the stand-alone exclusion index under expected shortfall at 99% is below the threshold $(v - \eta)/\kappa$. Panel (a): systematic standard deviation of the AI loss held by insurers, $c\,\mathrm{sd}(T)$, in units of a full-market event. Panel (b): loss of a lender whose borrowers mirror the market in a year in which the largest provider fails, $\ell\,\Psi(c)\,q \max_k \sigma_k$, and the increase in capital under the internal ratings-based formula at the re-estimated default probability, shown for the lower threshold, both in basis points of exposure. $p = 1\%$, $q = 10\%$, $r = 0.1$, PD $= 1\%$, LGD $= 40\%$, $\Psi_0 = 15\%$.

Timing. With 2023 shares the stand-alone index at $p = 1\%$ was 4.5, above its end-2025 value of 3.6, because the largest provider then held half of the market; the Herfindahl index was 0.29 in 2023 and 0.28 at the end of 2025, so the largest share fell while exclusions spread. The share of large EU enterprises using AI that generates language, by contrast, rose from 7.7% in 2023 to 31.7% in 2025, so that the share of large firms exposed to the largest provider rose from 3.9 to 12.7% if the shares of the US-centred surveys apply to European firms, and for a liability writer with equity of 60% of premium the scenario loss rose from $6.45z\%$ to $21.1z\%$ of equity, crossing an appetite of a tenth of equity for any z between 0.47 and 1.55 (the markers in Figure 6b). Filed adoptions of the standard exclusions began in August 2025 and peaked in February 2026, after this growth. With an illustrative appetite and intensity, and filings not matched to carriers' books, the timing is consistent with the appetite mechanism but does not test it.

Value-at-risk. Across value-at-risk at 99.5% and expected shortfall at 99 and 99.5%, stand-alone capital for end-2025 shares lies between 0.38 and 0.43 of a full-market event at p of 1 or 2%. The measures differ where Lemma 3 predicts. For a single-provider market with p of 0.25 or 0.4%, value-at-risk requires no capital, while expected shortfall at 99.5% requires 0.50 and 0.80, and at $p = 0.25\%$ value-at-risk assigns 0.24 to the end-2025 structure and nothing to a monopoly (Figure 5 and Table 16). Lemma 3 concerns the accumulation on its own. An internal model that adds the accumulation to the rest of the book before taking the quantile behaves differently, because a failure rarer than the tail probability still shifts the quantile of the total. If the rest of the book adds an independent normal loss whose standard deviation is a fraction s of a full-market event, the scale of other risks of Proposition 2 with standard normal U, and $p = 0.4\%$, the reversal of part (ii) between a single provider and ten equal providers survives only for s below about 0.06, and against seven equal providers below about 0.07. At $s = 0.1$, the increase in whole-book value-at-risk at 99.5% caused by the single-provider market, net of its expected loss, is 0.047 of a full-market event, against 0.574 for expected shortfall at 99.5%; at $s = 0.5$ the two are 0.035 and 0.061.

Computation. For a vector of shares σ and failure probabilities p, the distribution of T is obtained by enumerating the 2^K failure states and merging states with equal values of T. Value-at-risk is the lower α-quantile. Expected shortfall is computed exactly as $(1-\alpha)^{-1}\sum_j t_j \left|[C_{j-1}, C_j] \cap [\alpha, 1]\right|$, where t_j are the ordered support points and C_j the cumulative probabilities, which handles atoms correctly. Crossings of the exclusion frontier are found by root-finding on $\Xi(p)$, which is continuous in p. For integrated capital, the rest of the book is an independent normal loss with standard deviation s in units of a full-market event, $s = \mathrm{sd}_{\mathrm{rest}}/(a_w z)$; the quantile of the total is found by root-finding on the mixture distribution function, expected shortfall follows in closed form, and the Euler contribution of the accumulation is $\mathrm{E}[T \mid \text{total in the tail}] - \mathrm{E}[T]$, also in closed form. With a limit L per provider event, the accumulation is $T_L = \sum_k \min\{\sigma_k, L\} F_k$, whose distribution is enumerated in the same way as that of T. For the bank channel, the portfolio loss rate is $\ell\,[\mathrm{DR}(Z) + (1 - \mathrm{DR}(Z))\,\Delta(T)]$, where $\mathrm{DR}(Z)$ is the conditional default rate given the systematic factor Z (Vasicek 2002), $\Delta(T) = (\pi + qT)\Psi_0$ is the additional default probability from uninsured AI harms, and Z is independent of T; in the correlated variant, Z is shifted by one standard

deviation towards bad states in years in which any provider fails. Because the conditional default rate is a monotone function of Z, the distribution function of the portfolio loss rate is a finite mixture of normal distribution functions in closed form, and the 99.9% quantile is found by root-finding on it. Regulatory capital applies the Basel formula, without the maturity adjustment, to the default probability $\mathrm{PD} + (1 - \mathrm{PD})\,\mathrm{E}[\Delta(T)]$ that a bank would estimate from observed defaults. The counts in Table 11 are the numbers of hits returned by EDGAR full text search for each query and range of filing dates, retrieved on 27 October 2026; the share of filings referring to third-party AI that name OpenAI, ChatGPT or Anthropic is obtained by inclusion and exclusion from the counts of joint mentions. The code and data that reproduce every number, table and figure, including the hand-collected lists of provider events, insurance-market actions, contract terms and filing classifications, are deposited at Zenodo (Aslan 2026b).

Residual share. The split of the residual share among smaller providers has little effect at central values: with end-2025 shares and $p = 1\%$, expected shortfall capital at 99% is 0.397 whether the residual 12% is held by one provider or by twelve, because capital is driven by the largest provider. At $p = 0.25\%$ its effect is larger: across residuals held by one to twelve providers, capital ranges from 0.22 to 0.25 under expected shortfall at 99% and from 0.21 to 0.27 under value-at-risk at 99.5%.

Heterogeneous failure probabilities. Unequal failure probabilities change magnitudes. If the three large providers fail with probability 0.5% and the small ones with 2%, expected shortfall capital at 99% falls from 0.40 to 0.33; if the largest provider fails with probability 2%, the next two with 1% and the small ones with 0.5%, it is 0.40. They can also reverse the effect of concentration when larger providers are safer: two equal providers with failure probabilities of 2 and 0.2% need 0.49, while shares of 0.9 and 0.1 with the larger provider at 0.2% need 0.26.

Nested concentration. Shared infrastructure raises capital substantially. In a stylised example in which the providers run on two infrastructure platforms that host 67 and 33% of adopters' workloads and each fail with probability 0.5%, independently of provider failures with probability 1%, expected shortfall capital at 99% rises from 0.40 to 0.53 and at 99.5% from 0.40 to 0.66, and the systematic variance roughly doubles.

Capital concept. The concept of capital makes the largest difference: incremental capital for a diversified book is a small fraction of stand-alone capital at central values, and the Euler allocation lies in between (Table 17).

Default and intensity. The default probability of a harmed firm without cover is tied to the intensity through a lognormal loss per harm with mean z/q years of the firm's liability premium and coefficient of variation $\delta = 1$, and a buffer B beyond which the firm defaults. A buffer of 16.8 years of premium gives $\Psi_0 = 15\%$ at $z = 1$ and $q = 10\%$; at the same buffer, Ψ_0 is 3, 42 and 74% at $z = 0.5$, 2 and 4 (Table 18). At a given intensity, harms that are rarer and

larger cause more defaults: at $z = 1$, $q\Psi_0$ is 2.1% with $q = 5\%$ and 0.6% with $q = 20\%$. The four cases of Table 2 correspond to intensities between 0.5 and about 3 at this buffer.

Timing of losses. Liability claims are paid over several years. In the industry-aggregate Schedule P triangles of the CAS Loss Reserve Database for other liability on an occurrence basis, accident years 1988 to 1997 (Meyers and Shi 2011), 9% of an accident year's ultimate loss is paid in its first year and 21% in the third, the peak; for products liability the peak is 20% in the fourth year, and for medical malpractice on a claims-made basis 21% in the third. If defaults follow payments, the peak annual bank loss from a provider failure is therefore about a fifth of the failure-year figures reported in the text, which book the whole loss in one year.

Table 18. Failure-year losses with the default probability tied to the intensity

		Lender, bp of exposure		% of CET1 capital		
Intensity z	Ψ_0	Mirror	Concentrated	Euro-area SIs	Concentrated	Specialist
0.5	3%	5	12	0.07	0.2	1.0
1	15%	24	60	0.33	0.8	4.8
2	42%	67	168	0.93	2.3	13.4
4	74%	118	294	1.64	4.1	23.5

Notes: Failure of the largest provider ($\sigma_{\max} = 0.40$), $q = 10\%$, LGD $= 40\%$, default buffer of 16.8 years of liability premium. A mirror lender's adopting borrowers mirror the market; a concentrated lender's adopting borrowers all use the failing provider. Euro-area significant institutions: loans to non-financial corporations of EUR 6.2 trillion and CET1 capital of EUR 1,560 billion (European Central Bank 2026b), with 35% of corporate credit assumed to go to AI adopters. The concentrated lender holds the aggregate ratio of CET1 capital to corporate loans of euro-area significant institutions, about a quarter, and 35% of its corporate loans go to adopters; the specialist lender holds CET1 capital of an eighth of its loans, all to adopters of the failing provider. Losses are booked in the failure year.

Table 19. Bank losses on loans to AI adopters under full exclusion, corporate and SME portfolios

			Failure-year	Change under exclusion		
Portfolio	q	Ψ_0	extra loss	EL	Two-factor UL	IRB UL
Idiosyncratic and common harms ($\pi = 0.1q$)						
Corporate	5%	15%	12	3.3	−0.4	17.9
Corporate	10%	15%	24	6.5	−0.8	34.5
Corporate	10%	30%	48	13.1	−1.6	64.6
Corporate	20%	30%	96	26.1	−2.8	115.5
SME	5%	15%	12	3.2	−0.4	6.8
SME	10%	15%	24	6.5	−0.9	13.4
SME	10%	30%	48	12.9	−1.6	26.0
SME	20%	30%	96	25.9	−2.8	49.1
Provider failures alone ($\pi = 0$)						
Corporate	5%	15%	12	0.3	0.0	1.7
Corporate	10%	15%	24	0.6	0.0	3.3
Corporate	10%	30%	48	1.2	0.0	6.6
Corporate	20%	30%	96	2.4	0.4	13.1
SME	5%	15%	12	0.3	0.0	0.6
SME	10%	15%	24	0.6	0.0	1.3
SME	10%	30%	48	1.2	0.0	2.5
SME	20%	30%	96	2.4	0.4	5.0

Notes: As Table 2, with SME portfolios added. PD of 1% for corporate and 2% for SME portfolios; the asset correlation follows the Basel function, with the size adjustment at EUR 10 million of turnover for SMEs, and is re-evaluated at the new default probability. All figures assume full exclusion and scale with the share of borrowers whose AI-related liability is excluded.

Table 20. Failure-year loss across banks in the EBA 2025 transparency exercise

	10th pct.	Median	90th pct.	Maximum
Share of NFC loans to sections J and M, %	3	8	15	57
Sector-weighted adoption, relative to all banks	0.83	0.97	1.11	1.76
Loans to NFCs relative to CET1 capital	1.43	3.63	6.26	10.5
Failure-year loss, % of CET1 capital, $q = 10\%$, $\Psi_0 = 15\%$				
Borrowers mirror the market	0.10	0.28	0.53	0.87
Borrowers in J and M on the failing provider	0.12	0.37	0.66	1.03
Adopting borrowers all on the failing provider	0.25	0.70	1.33	2.17
Same, % of headroom above CET1 requirements				
Borrowers mirror the market	0.25	0.83	2.22	4.69
Adopting borrowers all on the failing provider	0.63	2.07	5.56	11.7

Notes: 104 banks from 25 EU and EEA countries, together with one bank whose country is not disclosed, 30 June 2025 (European Banking Authority 2025). Four further banks that report these loans are excluded because their risk-weighted assets are less than a tenth of their total assets, against 36% at the median bank, which indicates lending to or guaranteed by the public sector; with a loss given default of 40%, they would lose 0.22 to 1.57% of CET1 capital if their borrowers mirrored the market and 0.55 to 3.93% if all adopting borrowers relied on the failing provider. Headroom is the CET1 ratio less 11.2%, the average CET1 requirements and guidance of significant institutions for 2026 (European Central Bank 2025), applied to all banks. Loans to NFCs are gross carrying amounts of loans and advances to non-financial corporations other than held for trading, by NACE section; J is information and communication, M professional, scientific and technical activities. Sector-weighted adoption is the loan-weighted share of enterprises with ten or more persons employed that use AI generating written or spoken language in 2025 (Eurostat 2026b), relative to the same average for all banks; sections outside the Eurostat survey (agriculture, mining, financial activities and sections O to S) take the rate of all surveyed activities. The failure-year loss is NFC loans $\times$ 0.35 $\times$ relative adoption $\times$ LGD $\times\, q \times$ share $\times\, \Psi_0$ over CET1 capital, where the share of adopting borrowers on the failing provider is 0.40 when borrowers mirror the market and 1 when all adopting borrowers rely on it; in the row for sections J and M, borrowers in those sections rely on the failing provider and the others mirror the market; LGD $= 40\%$. Percentiles are computed separately for each row, so a column does not describe a single bank.

D Provider-level events and insurance-market actions

This appendix summarises and lists the provider-level events of Section 2.2 (Tables 21 and 24) and the insurance-market actions of Section 2.4 (Tables 22, 23 and 25).

Table 21. Documented provider-level events at eight foundation-model providers

Type of event	2023	2024	2025	2026	Total
Behaviour changes and defects	2	4	9	1	16
Regressions in safeguards	0	0	4	0	4
Outages of two hours or more	1	4	5	7	17
Suspensions and withdrawals	1	1	2	4	8
Security incidents	2	2	3	2	9
Other	0	0	0	1	1
All events	6	11	23	15	55

Notes: Events at OpenAI, Anthropic, Google, Microsoft, Amazon Web Services, Meta, xAI and Mistral in which one change or failure at the provider reached many of its clients at once, compiled from providers' postmortems, status pages and press reports; 2026 runs to 28 October. Outages are counted if they lasted about two hours or more and affected most customers of a service, or all customers of a service in a region. Status pages log many shorter and partial incidents; the full list is in Table 24 and the coding rules below.

Table 22. Documented insurance-market responses to AI liability, January 2023 to October 2026

	Multi-line	Specialty	Lloyd's or MGA	Insur-tech	Rein-surer	Bro-ker	Other	Total
Exclusions and sublimits								
AI in general	2	2	–	–	–	–	–	4
Generative AI	–	3	–	–	–	–	2	5
Widespread simultaneous losses	1	–	–	–	–	–	–	1
Sublimit for AI losses	2	1	–	–	–	–	–	3
Affirmative products								
Liability cover for AI users	–	2	7	1	–	–	–	10
Cyber cover for AI events	–	2	1	3	–	–	1	7
AI developers and vendors	1	1	–	4	–	1	–	7
Performance cover	–	1	2	–	2	–	–	5
Statements								
Accumulation or capacity	–	1	1	–	2	6	2	12
All actions	6	13	11	8	4	7	5	54

Notes: Number of documented actions by content and type of company. Each published form, filing, wording, product or statement counts as one action, so a company can account for more than one: the 13 restrictions come from ten named companies, two of which (W. R. Berkley and QBE) account for two each, and from the unnamed carriers in the market-wide count of filings. Exclusions and sublimits include filings with state regulators; exclusions of generative AI include the standard ISO forms (entry 6 in Table 25) and the market-wide count of their adoption (entry 18), each entered once. Liability cover for AI users is third-party cover for firms that deploy AI; cyber cover for AI events includes endorsements for AI-enabled attacks and AI system failures in cyber policies; performance cover pays when a model misses defined performance thresholds. Statements are public statements by carriers, reinsurers, brokers or market bodies that tie AI cover to accumulation or capacity. "Other" comprises the form provider, the market-wide count of filings, Lloyd's as a market and an insurance programme of a cloud provider. Ten of the actions counted rest on a single report or on a second-hand account; Table 25 marks them, and also lists, without counting them here, filings by five named carriers that the market-wide count already includes.

Table 23. What three explanations of AI exclusions predict, and what the documented actions show

Explanation	Predicted wordings, limits and placement	Documented
Contract certainty, as with silent cyber	Exclusions of AI or generative AI as such, following standard forms and uniform across insureds; cover sold back by endorsement	9 of 13 restrictions; 4,078 adoptions of the standard forms; 10 liability products for firms that use AI and 7 cyber endorsements
Unobservable use of AI by the insured	Cover conditioned on audits, testing or verification of the insured's systems; exclusions varying with observable AI intensity	Audit-linked and verification-linked products of two firms
Accumulation	Exclusions at carriers whose books are dominated by AI-exposed insureds; wordings for simultaneous losses; limits per event or in aggregate; capacity scarcest for the largest insureds	1 exclusion of widespread simultaneous losses; 1 draft sublimit for specified AI losses; limits per insured only; no aggregate per provider event; 12 statements on accumulation or capacity

Notes: Counts of documented actions from Table 22; a company can account for more than one action. Two further sublimits, for fines under the EU AI Act and, in cyber wordings, for the use of stolen credentials to run AI models, concern particular perils and are attributed to none of the explanations. When a carrier cannot observe which provider an insured depends on, an exclusion of AI as such is also what the accumulation predicts (Section 3.5), so the counts in the first row do not discriminate between the first and third explanations.

D.1 Provider-level events

Table 24 lists the events behind Table 21, with the site of each source, linked in the electronic version to the source itself. They were compiled from providers' postmortems and status pages and from press reports, for OpenAI, Anthropic, Google, Microsoft, Amazon Web Services, Meta, xAI and Mistral, from January 2023 to 28 October 2026. The coding rules are as follows.

(i) An event is counted if one change or failure at the provider reached many of its clients at once.

(ii) *Behaviour* covers changes in what a model produces that the provider rolled back or acknowledged as defects; *Safeguards*, regressions in safety measures; *Outage*, unavailability of about two hours or more affecting most customers of a service, or all customers of a service in a region; *Suspension*, regulatory orders, regional restrictions and the retirement of a model or product for all of its users, including announced retirements (events 47 and 49); *Security*, exposure of customers' data or compromise of a component.

(iii) Routine deprecations of interface versions with a successor version, shorter incidents and partial incidents that affect particular models or some customers within a region are not counted (event 13, a series of seven partial incidents). An outage that a status page labels partial but that affects most customers of a service, as event 55 did, is counted. Counted events that turn on these rules are the three single-region outages (events 39, 45 and 57) and the two announced retirements (events 47 and 49); without them the count is 50. Event 19 is counted because it affected the Gemini 1.5 models, then the generally available models of Google's interface, in 29 locations. Changes that the provider did not acknowledge as defects are not counted (event 24, the remapping of a dated interface endpoint to a different model).

(iv) An event *reached interface customers* (marked $*$) if the provider's own account states

that the change affected requests made through its interface, including through cloud marketplaces. Events 6 and 32 qualify. Event 21, the uneven quality of a new family of open-weight models, arose in deployments by third-party hosts, and event 48 did not affect the interface.

The two events marked $*$ occurred at OpenAI (event 6) and Anthropic (event 32), two of the four providers with the largest shares in the surveys; the others are Google and Meta. Meta's open-weight models reached most users through other firms' hosting services, and Meta offered its own interface only from April 2025 (Wiggers 2025). Counting Meta's provider-years from then gives 13.0 provider-years for the four providers from January 2023 to October 2026 and a rate of 0.15 a year, with an exact Poisson 95% interval from 0.019 to 0.56; counting them from January 2023 gives 15.3 provider-years and a rate of 0.13 a year, with an interval from 0.016 to 0.47. If between one in a hundred and one in ten such events produced harm at scale, the annual probability of a liability-relevant failure per provider would lie between 0.016% and 5.6%, the lower end from the interval with 15.3 provider-years and the upper end from that with 13.0, a range that contains the grid of Appendix C.1. Microsoft's cloud service is not counted as a separate provider because it serves OpenAI's models, whose defects are counted under OpenAI.

Table 24. Documented provider-level events, January 2023 to October 2026

No.	Date	Provider	Type	Scope and duration	Counted	Source
1	2023-02-17	Microsoft	Behaviour	Bing Chat (consumer); provider said long sessions could confuse the model; session limits from the second week of the preview	Y	cnbc.com
2	2023-03-20	OpenAI	Security	ChatGPT chat titles and Plus payment details; 9 h exposure window	Y	openai.com
3	2023-03	OpenAI	Suspension	ChatGPT in Italy (regulator order); about 4 weeks	Y	fortune.com
4	2023-09-28	Google	Security	Bard shared chats indexed by search	Y	theregister.com
5	2023-11-08	OpenAI	Outage	API and ChatGPT (DDoS); about 25 h of periodic outages	Y	status.openai.com
6	2023-12-08	OpenAI	Behaviour	ChatGPT GPT-4 and API GPT-4 Turbo preview; at least 48 days	Y*	openai.com
7	2024-01-22	Anthropic	Security	subset of customers' account names and balances; one-off	Y	venturebeat.com
8	2024-02-20	OpenAI	Behaviour	ChatGPT; about 7.5 h	Y	status.openai.com
9	2024-02-22	Google	Behaviour	Gemini app image generation of people; about 6 months (feature paused)	Y	techcrunch.com
10	2024-05	Google	Behaviour	Search AI Overviews (consumer); about 2 weeks	Y	wabe.org
11	2024-07-09	OpenAI	Suspension	API traffic from unsupported regions; permanent	Y	theregister.com
12	2024-07-13	Microsoft	Outage	Azure OpenAI in 14 of 28 regions; 23.6 h (64 h for fine-tuned deployments)	Y	azure.status.microsoft
13	2024-11-08	Google	Outage	Vertex Gemini API specific models or regions; 2.7 h to 4.8 days (7 partial incidents)	N	status.cloud.google.com
14	2024-11-25	OpenAI	Outage	ChatGPT all tiers and some API models; about 2.75 h (API)	Y	status.openai.com
15	2024-12-11	OpenAI	Outage	all services; 4.4 h	Y	status.openai.com
16	2024-12-26	OpenAI	Outage	many API endpoints and ChatGPT (cloud data-centre power failure); 4.5 h (API)	Y	status.openai.com
17	2024-12-26	Meta	Security	Meta AI prompts visible to other users; 29 days	Y	techcrunch.com
18	2024-12	Microsoft	Behaviour	Bing Image Creator model rollback; several weeks	Y	techcrunch.com

Table 24, continued

No.	Date	Provider	Type	Scope and duration	Counted	Source
19	2025-01-07	Google	Outage	Vertex Gemini 1.5 models in 29 locations; 7.8 h	Y	status.cloud.google.com
20	2025-02-23	xAI	Behaviour	Grok 3 system prompt; under 1 day	Y	windowscentral.com
21	2025-04-05	Meta	Behaviour	Llama 4 across third-party hosts; several days	Y	venturebeat.com
22	2025-04-25	OpenAI	Behaviour	GPT-4o update in ChatGPT (sycophancy); about 4 days	Y	openai.com
23	2025-04-28	OpenAI	Safeguards	ChatGPT accounts of minors	Y	techcrunch.com
24	2025-05-08	Google	Behaviour	dated Gemini API endpoint remapped; not determined	N	discuss.ai.google.dev
25	2025-05-14	xAI	Behaviour	Grok replies on X (prompt modification); hours	Y	tech.yahoo.com
26	2025-06-09	OpenAI	Outage	API (error rates up to about 25%) and ChatGPT; 15.4 h	Y	status.openai.com
27	2025-06-12	Google	Outage	Google Cloud incl. Vertex Gemini API; 7.5 h	Y	status.cloud.google.com
28	2025-07-07	xAI	Behaviour	Grok replies on X (upstream code path); 16 h	Y	gigazine.net
29	2025-07-15	xAI	Behaviour	Grok 4 responses; days	Y	techcrunch.com
30	2025-07-17	AWS	Security	Amazon Q Developer extension; 7 days	Y	aws.amazon.com
31	2025-07-31	OpenAI	Security	shared chats indexed by search engines; 1 to 2 days	Y	theregister.com
32	2025-08-05	Anthropic	Behaviour	Claude API incl. Bedrock and Vertex (routing bug); 8 to 44 days (three bugs)	Y*	anthropic.com
33	2025-08-07	OpenAI	Behaviour	ChatGPT model switch at GPT-5 launch; about 4 days	Y	techrepublic.com
34	2025-08-08	Google	Behaviour	Gemini looping responses	Y	androidauthority.com
35	2025-08-14	Meta	Safeguards	Meta AI chatbot guidelines	Y	cp24.com
36	2025-08-26	OpenAI	Safeguards	safety training degrading in long ChatGPT conversations; ongoing	Y	openai.com
37	2025-09-04	Anthropic	Suspension	firms majority-owned from unsupported regions; permanent	Y	anthropic.com
38	2025-10-03	Anthropic	Outage	claude.ai Console API Claude Code; 2.5 h to mitigation	Y	status.claude.com
39	2025-10-19	AWS	Outage	us-east-1 many services; 14.5 h	Y	aws.amazon.com
40	2025-11-01	Google	Suspension	Gemma removed from AI Studio; permanent	Y	techcrunch.com
41	2025-11-09	OpenAI	Security	API account profile data via analytics vendor	Y	openai.com
42	2025-12	xAI	Safeguards	Grok editing images of real people; 2 to 3 weeks	Y	decrypt.co
43	2026-01-10	xAI	Suspension	Grok blocked in three countries; up to about 3 weeks	Y	deccanherald.com
44	2026-01-21	Microsoft	Security	Copilot Chat summarising labelled emails; about 1 month	Y	bleepingcomputer.com
45	2026-01-27	Microsoft	Outage	Azure OpenAI Sweden Central; 7.2 h	Y	theregister.com
46	2026-02-03	OpenAI	Outage	ChatGPT and platform login; 2.9 h	Y	status.openai.com
47	2026-02-13	OpenAI	Suspension	older models retired in ChatGPT; permanent (scheduled)	Y	openai.com
48	2026-03-04	Anthropic	Behaviour	Claude Code Agent SDK and Cowork (API not affected); 47 days	Y	anthropic.com
49	2026-04-26	OpenAI	Suspension	Sora app and API; permanent (scheduled)	Y	help.openai.com
50	2026-05-11	Mistral	Security	compromised SDK packages; about 3 h per registry	Y	docs.mistral.ai
51	2026-06-02	Anthropic	Outage	claude.ai Console API Claude Code; 5.75 h	Y	status.claude.com
52	2026-06-08	Microsoft	Outage	Azure OpenAI in several European regions; 5.5 to 7.25 h	Y	status.templafy.com
53	2026-06-12	Anthropic	Suspension	two frontier models suspended for all customers (export control); 19 days	Y	anthropic.com
54	2026-07-25	OpenAI	Outage	ChatGPT API Platform Codex; about 3 h	Y	status.openai.com
55	2026-09-03	Anthropic	Outage	partial outage incl. API (outages at two other providers the same morning); 3.1 h	Y	theregister.com
56	2026-09-18	OpenAI	Other	billing defect in Agent API; 9.5 h	Y	status.openai.com
57	2026-09-29	Microsoft	Outage	Azure OpenAI and Foundry Sweden Central; 5.9 h	Y	azure.status.microsoft

Table 24, continued

No.	Date	Provider	Type	Scope and duration	Counted	Source

D.2 Insurance-market actions

Table 25 lists the actions behind Table 22, again with the site of each source. They were compiled from company announcements and policy forms, from filings reported in the trade press and from broker and reinsurer reports, for January 2023 to October 2026, and each entry rests on a source page, linked in the electronic version; the replication files give a supporting passage for each. Restrictions are exclusions, sublimits and their filings with state regulators; affirmative products include endorsements that grant AI cover explicitly; statements tie AI cover to accumulation or to the capacity of the market. The column "Content" gives the coding used in Table 22: exclusions of AI in general (AI), of generative AI (GenAI) or of widespread simultaneous losses (Widespread), and sublimits; liability cover for firms that use AI (Deployer), cyber cover extended to AI events (Cyber), cover for AI developers and vendors (Vendor) and performance cover and warranties (Performance); and statements. The coding follows the wording of the source; products that serve both deployers and developers are coded by the first group the source names. Where the source does not date the action, the date is that of the report. Entries marked † rest on a single report or on a second-hand account of another outlet's reporting, such as reports that cite the Financial Times. Entries marked ‡ are filings by named carriers that the market-wide count of filings (entry 18) already includes; Table 22 does not count them separately. Three entries that reported only the funding of a firm, without a product or its terms, or a filing that the insurer named has denied planning to use, are left out; entry 43, which reports a funding round together with a product and its limits, is counted as a product.

Table 25. Documented insurance-market actions on AI liability, January 2023 to October 2026

No.	Date	Company	Type	Content	Action	Limit	Source
Exclusions and sublimits							
1	2024-06	W. R. Berkley (Berkley Insurance Company)	Specialty	AI	Absolute AI exclusion endorsement (PC 51380) for management-liability cover, barring claims from any use of AI		hunton.sitepilot11.firmseek.com
2	2025-10-28	W. R. Berkley	Specialty	AI	Filed one of the broadest definition-based AI exclusions with US state regulators		theinsurer.com
3	2025-10-28	Cincinnati Insurance (Cincinnati Financial)	Multi-line	AI	Filed an AI exclusion that defines AI by combining the OECD and White House definitions		theinsurer.com
4	2025-10-28	Frederick Mutual Insurance	Multi-line	AI	Filed an AI exclusion that reaches beyond generative tools		theinsurer.com
5	2025-10-28	Philadelphia Insurance Companies (Tokio Marine)	Specialty	GenAI	Filed a generative-AI exclusion that defines generative AI by the content it produces		theinsurer.com

Table 25, continued

No.	Date	Company	Type	Content	Action	Limit	Source
6	2025-07	Verisk (ISO)	Other	GenAI	ISO generative-AI exclusions CG 40 47, CG 40 48 and CG 35 08 (products), effective January 2026		independentagent.com
7	2026-01-22	Chubb†	Multi-line	Wide-spread	Covers some AI-related losses but excludes incidents causing widespread simultaneous losses		roughnotes.com
8	2026-01-22	QBE†	Multi-line	Sublimit	Sublimit of up to 2.5% of the insured limit for fines under the EU AI Act	2.5% of insured limit	roughnotes.com
9	2026-02-01	Great American Insurance Group†	Specialty	GenAI	Amended Washington umbrella and excess forms to exclude generative-AI liability		beinsure.com
10	2026-04	Travelers†‡	Multi-line	GenAI	Filed to adopt the ISO generative-AI exclusions or equivalent wordings by April 2026		insurancebusiness-mag.com
11	2026-04	Berkshire Hathaway†‡	Multi-line	GenAI	Filed to adopt the ISO generative-AI exclusions or equivalent wordings by April 2026		insurancebusiness-mag.com
12	2026-04	Chubb†‡	Multi-line	GenAI	Filed to adopt the ISO generative-AI exclusions or equivalent wordings by April 2026		insurancebusiness-mag.com
13	2026-04-14	Hamilton Insurance Group†	Specialty	GenAI	Exclusion of claims arising out of or in any way involving generative AI tools		jdsupra.com
14	2026-04-21	Liberty Mutual‡	Multi-line	GenAI	One of seven insurers with more than 60% of generative-AI exclusion filings, January to April 2026		independentagent.com
15	2026-04-21	AXA (US filings)‡	Multi-line	GenAI	One of seven insurers with more than 60% of generative-AI exclusion filings, January to April 2026		independentagent.com
16	2026-04-23	QBE†	Multi-line	Sublimit	Sublimits for LLM-jacking incidents in cyber wordings; the insurer says it is not retreating from AI risk	about $250,000 on a $5m cyber policy	finance.yahoo.com
17	2026-04-23	Beazley†	Specialty	Sublimit	Draft terms capping specified AI losses at about 10% of the policy limit, not yet in any policy	about 10% of overall policy value	finance.yahoo.com
18	2026-07-31	Multiple US carriers (unnamed; Trades Coverage filing analysis)	Other	GenAI	4,078 filed adoptions of the ISO generative-AI exclusions in 49 states and DC by 31 July 2026		insurancebusiness-mag.com
Affirmative AI products							
19	2023-04-28	Munich Re	Reinsurer	Perform.	Performance cover (aiSelf) for firms using self-developed AI and machine-learning models		reinsurancene.ws
20	2023-09-28	Armilla Assurance (Armilla AI)	Lloyd's or MGA	Perform.	Performance warranty for verified AI products that refunds licence fees if targets are missed		armilla.ai

Table 25, continued

No.	Date	Company	Type	Content	Action	Limit	Source
21	2024-01-30	Vouch	Insurtech	Vendor	Insurance for AI start-ups: errors and omissions, algorithmic bias, IP claims and regulatory defence		insurance-canada.ca
22	2024-03-26	Coalition	Insurtech	Cyber	Affirmative AI endorsement for cyber policies, covering AI security events and deepfake fraud		coalitioninc.com
23	2024-10-21	AXA XL	Multi-line	Vendor	Generative-AI endorsement to cyber cover for firms that develop their own models		axaxl.com
24	2025	Coalition	Insurtech	Cyber	Affirmative AI endorsement made part of the base cyber policy		businessinsurance.com
25	2025-01-14	Relm Insurance	Specialty	Vendor	Three products for AI firms and users: tech E&O, an excess wrap where AI is excluded, first-party response		relminsurance.com
26	2025-04-30	Armilla Insurance Services (Armilla AI)	Lloyd's or MGA	Deployer	Affirmative AI liability insurance triggered by AI errors, hallucinations and underperformance		armilla.ai
27	2025-05-14	Hiscox (Hiscox Underwriting Ltd, UK)	Specialty	Deployer	UK technology professional indemnity wording with explicit cover for firms that use or build AI		hiscox.co.uk
28	2025-05-15	Google Cloud Risk Protection Program (insurer partners Beazley, Chubb, Munich Re)	Other	Cyber	Affirmative cover for customers' AI workloads in a cloud provider's insurance programme		cloud.google.com
29	2025-07-23	Artificial Intelligence Underwriting Company (AIUC)	Insurtech	Vendor	Agent standard, testing and liability cover for AI vendors and their customers when an agent fails		reinsurancene.ws
30	2025-10-08	Aon (broker; placement for OpenAI)†	Broker	Vendor	OpenAI reportedly arranged up to USD 300m of cover; another source put the figure much lower	up to $300m	businessinsurance.com
31	2025-11-26	Counterpart	Insurtech	Deployer	Affirmative AI cover in professional liability and a new technology E&O insuring agreement		fintech.global
32	2026-01	Testudo (Testudo Global)	Lloyd's or MGA	Deployer	Third-party generative-AI liability insurance for enterprises that deploy generative AI		iua.co.uk
33	2026-01-21	Armilla AI	Lloyd's or MGA	Deployer	Higher stand-alone AI liability limits and cover extended to general liability	up to $25m per organization	armilla.ai

Table 25, continued

No.	Date	Company	Type	Content	Action	Limit	Source
34	2026-02-10	Chaucer and Armilla AI (Vanguard AI)	Lloyd's or MGA	Deployer	Cyber and tech E&O cover combined with a stand-alone AI liability policy	AI aggregate $25m or more per organisation; $10m cyber	chaucergroup.com
35	2026-02-26	Mosaic Insurance (with Munich Re aiSure)	Specialty	Perform.	Performance cover that pays when a model misses defined thresholds, without proof of negligence	up to $15m	businessinsurance.com
36	2026-03-09	Testudo	Lloyd's or MGA	Deployer	Added two carriers to a generative-AI liability panel and raised the available limits	$9.25m per insured	atrium-uw.com
37	2026-03-18	HSB (Hartford Steam Boiler; Munich Re)	Specialty	Deployer	AI liability cover for small and mid-sized firms, added to partner carriers' business policies		munichre.com
38	2026-03-18	iTOO Special Risks (with Munich Re)	Lloyd's or MGA	Perform.	Cover in South Africa for non-malicious failure of AI systems to perform as intended		it-online.co.za
39	2026-04-21	Cowbell	Insurtech	Cyber	Cyber product with affirmative cover for AI-related incidents, for firms with USD 250m to 1bn of revenue	up to $10m	cowbell.insure
40	2026-05-06	Corgi Insurance	Insurtech	Vendor	Modular AI cover within tech E&O for biased outputs, training-data misuse and adversarial attacks		fintech.global
41	2026-06-24	Mayflower Specialty (MGA) with Hadron	Lloyd's or MGA	Deployer	Affirmative AI liability programme across D&O, employment practices and E&O, with an excess layer		businesswire.com
42	2026-06-26	CFC	Lloyd's or MGA	Deployer	Affirmative AI wording in seven core products, covering hallucination and model drift		reinsurancene.ws
43	2026-09-15	Artificial Intelligence Underwriting Company (AIUC)†	Insurtech	Vendor	Audit-linked insurance for AI agents, reportedly extended to frontier models; announced with Series A funding	up to $50m per policy	dealroom.co
44	2026-09-17	Beazley	Specialty	Cyber	Affirmative AI endorsement in cyber and tech E&O policies		insurancejournal.com
45	2026-09-17	CFC	Lloyd's or MGA	Cyber	Affirmative cover for AI-related cyber exposures in a product for financial institutions		insurancejournal.com
46	2026-09-25	Beazley	Specialty	Cyber	Endorsements for shutting down a malfunctioning AI system and for AI regulatory defence and penalties		insurancebusinessmag.com
47	2026-10-16	Munich Re (aiSure)†	Reinsurer	Perform.	Cover for errors of AI models with limits of up to USD 50m	up to $50m	itiger.com
Statements on accumulation or capacity							
48	2023-11-03	Beazley	Specialty	Statement	Chief underwriting officer: an AI disaster scenario will set the appetite, and exclusions are possible		businessinsurance.com

Table 25, continued

No.	Date	Company	Type	Content	Action	Limit	Source
49	2024-03	Lloyd's of London (Futureset report)	Other	Statement	Report: concentration among model providers creates single points of failure and large accumulations		assets.lloyds.com
50	2024-03-21	Munich Re	Reinsurer	Statement	Head of AI insurance: one model used across many companies creates systematic accumulation risk		reinsurancene.ws
51	2025-10-08	Aon	Broker	Statement	Head of cyber risk: the market lacks capacity for AI model providers		insurancejournal.com
52	2025-11-23	Aon	Broker	Statement	Executive: a USD 400m loss at one company is insurable, 10,000 simultaneous losses are not		techcrunch.com
53	2026-02-02	Swiss Re (Swiss Re Institute)	Reinsurer	Statement	Report: reliance on a few cloud and AI providers raises systemic and concentration risk		insurancenews.com.au
54	2026-02-05	Lockton Re (with Armilla AI)	Broker	Statement	Report: systemic risk from shared AI infrastructure and common foundation models		global.lockton.com
55	2026-03-24	Gallagher Re	Broker	Statement	White paper: faults in widely adopted AI systems could trigger simultaneous claims across sectors		reinsurancene.ws
56	2026-06-10	Gallagher Re	Broker	Statement	Report: one flaw in a shared model can hit many insureds at once, and insurers respond by excluding AI		ajg.com
57	2026-08-27	Verisk (Verisk Underwriting Solutions)	Other	Statement	Analytics firm: systemic events in which one AI model causes losses at many organisations		businessinsurance.com
58	2026-10-15	Marsh (Marsh McLennan)	Broker	Statement	Broker's chief executive, opposing broad exclusions: understand how AI risks aggregate first		insurancejournal.com
59	2026-10-16	Testudo[†]	Lloyd's or MGA	Statement	Chief executive: cover at the scale the largest AI companies need does not exist		itiger.com

E Equity prices around provider incidents

On the afternoon of Friday 17 November 2023, OpenAI's board removed its chief executive (OpenAI 2023). Late on Sunday 19 November, Microsoft announced that he and colleagues would join it to lead a new research team, having stated that its long-term agreement with OpenAI gave it full access to everything it needed (Microsoft 2023), and most of OpenAI's staff threatened to follow him. On the evening of Tuesday 21 November, after the close of trading, OpenAI announced an agreement in principle for his return (Semafor 2023). The models did not change, but for three trading days the continuity of a provider on which many firms had built products was in doubt, until an announcement made after the close on the third day settled it.

Sample. I identify deployers from their own filings with the US Securities and Exchange Commission. A full-text search of current, quarterly, annual and foreign issuers' reports filed between 1 January and 16 November 2023 that mention OpenAI, together with filings that use the phrases "ChatGPT plugin" or "OpenAI's GPT", returns 66 firms; the filings retained were filed between 22 February and 14 November 2023, 40 of them on Form 8-K and 17 on Form 6-K. Each firm is classified from one filing, read in full for 57 of the 66 firms, by what it says about OpenAI. Forty-one firms state that they use OpenAI's models or interfaces in their own products, services or operations, seven of them describing a product as built on those models; 17 mention OpenAI or ChatGPT only as an example, in risk factors or in commentary on the market; two partners, two competitors that develop their own models and four firms whose filings could not be classified are set aside. The classification, the passages on which it rests and the links to the filings are in the replication files. The sample keeps firms with a share price of at least USD 5 on 16 November 2023 whose price history is available: 21 deployers and 8 firms with a generic mention. Of the 41 deployers, 9 had a share price below USD 5 and 11 have no price history that could be retrieved; of the 17 firms with a generic mention, 6 and 3.

Design. Returns are computed from closing prices adjusted for splits and dividends for 1 November to 1 December 2023, read from the historical price pages of Yahoo Finance (Yahoo Finance 2026); for the price filter, closing prices were converted to the prices quoted at the time using the splits reported after 16 November 2023, which changed the classification of three firms whose reverse splits had lifted their adjusted prices above USD 5, and a second retrieval of three firms' prices over a different range reproduced every value. Abnormal returns subtract the return of an S&P 500 index fund (SPY) or of a software-sector index fund (IGV). Inference uses the cross-section of firms and the other windows of the same month. The design fixed before the returns were examined, a market model estimated over the preceding year, could not be implemented because full price histories could not be retrieved, so the market-adjusted design reported here departs from it; the script for the pre-specified design is in the replication files and would also allow the sample to be enlarged to firms that disclose dependence only in annual reports.

Results. Table 26 and Figure 9 show no robust difference on average. Relative to the software index, deployers lost 1.2% during the crisis and gained 1.1% when it was resolved, within the range of their own returns in other three-day windows that month (-2.2 to 3.8%); relative to the S&P 500 their crisis return was nil. The difference between the groups in other three-day windows ranged from -1.2 to 2.0 points. The difference from firms that only mention OpenAI in the crisis window was -0.4 percentage points, with a permutation p-value of 0.74 and a 95% confidence interval from -3.1 to 2.4 points. With a standard error of 1.25 points, the test could detect only differences of about 3.5 points or more. The comparison group is itself imperfect, since five of its eight firms have some technical link with OpenAI's models: without Thomson Reuters, which was acquiring a product built on them, the difference is -0.3 points, and without all five it is -1.1 points, against three remaining firms. On the day of resolution the difference, 2.9 points, exceeds that of any other day that month (-1.2 to 1.2), but it comes from one firm in the comparison group whose shares fell by a fifth that day; without it the

Table 26. Share-price reactions to OpenAI's governance crisis, 17–22 November 2023

	n	Crisis, 17–21 Nov.	Resolution, 22 Nov.
Market-adjusted returns (%)			
Deployers	21	0.00	0.71
of which built on OpenAI's models	4	0.01	0.10
Generic mention	8	0.36	−2.21
Returns adjusted by the software index (%)			
Deployers	21	−1.15	1.12
of which built on OpenAI's models	4	−1.13	0.51
Generic mention	8	−0.78	−1.80
Deployers minus generic mention			
Difference (percentage points)		−0.37	2.92
Welch p-value		0.78	0.33
Permutation p-value		0.74	0.10
Rank-sum p-value		1.00	0.18
95% confidence interval		−3.13 to 2.40	
Placebo days or windows in November, range		−1.16 to 1.97	−1.16 to 1.18

Notes: Equal-weighted mean abnormal returns. Crisis: cumulative abnormal return on 17, 20 and 21 November 2023; resolution: abnormal return on 22 November. Deployers state in filings made before the crisis that they use OpenAI's models in their products, services or operations; generic mention: firms that mention OpenAI or ChatGPT only as an example. Market-adjusted returns subtract the return of SPY, and the second panel that of the software-sector fund IGV; the difference between the groups is the same under both. Permutation p-values from 100,000 random reassignments of the two labels; rank-sum: Mann–Whitney test. Confidence interval from Welch's t distribution. Placebo: the difference between the groups in every other window of three consecutive trading days (crisis) or every other day (resolution) in November 2023 outside 17–22 November.

difference is 0.2 points. The firms lost to the price filter and to later delistings are mostly small ones, which tilts the sample towards large deployers with diversified businesses and the test towards a null. The design that dependence data allow is a comparison of firms by their disclosed provider, which Section 6.3 implements across the incidents since 2023 for firms that disclose dependence in annual reports.

Two contested incidents. Two later events are poor tests taken alone. The suspension of two models in June 2026 (event 53 in Table 24) left the provider's other models available to all customers (Anthropic 2026b), so clients could switch at once, and a suspension makes a model unavailable without changing what it produces. The defects of August and September 2025 (event 32) did change outputs, but they developed over several weeks and were described in full only in a postmortem published after they had been fixed (Anthropic 2025a), so the date on which news of them reached investors is uncertain. Both are among the 38 incidents of Section 6.3; without them the difference over days 0 and 1 is −1.36 percentage points ($t = -3.3$), and without all suspensions it is −1.94 ($t = -5.2$).

Sample of the cross-provider test. The firms are those coded in Section 6.1 as selling products that run on a provider's models, with the providers they name. The coding rests on the passages kept from each annual report; a firm is coded as dependent when it states that a product or service it sells uses, integrates, is built on or is powered by a named provider's model, and as an internal user when the use it describes is confined to its own staff or operations. Connectors that let a firm's customers call a model of their choice, applications



Figure 9. Cumulative abnormal returns around OpenAI's governance crisis
Notes: Equal-weighted portfolios of deployers (21 firms) and of firms with a generic mention (8), returns adjusted by the software-sector fund IGV, cumulated from 16 November 2023. The shaded area marks the crisis window, 17 to 21 November.

distributed inside a provider's chatbot and mentions in lists of competitors or examples are not coded as dependence. Seven ticker matches to companies on other exchanges, found on reading the matched names, were discarded. Prices are adjusted daily closes from January 2022 to 11 November 2026 for 80 of the 89 firms, the S&P 500 index fund (SPY) and the software-sector fund (IGV) (Yahoo Finance 2026); daily returns above 80% in absolute value are treated as missing. A firm enters a test as a client of a provider from twelve months before the filing date of its first annual report that names it, since an annual report describes the preceding fiscal year.

Incidents. The incidents are the counted events of Table 24 at OpenAI, Microsoft (whose Azure OpenAI service is counted under OpenAI), Google, Anthropic, Meta and Mistral with a precise date; events at xAI and Amazon Web Services are not used, since no coded firm names xAI's models and only two name Amazon's hosting service, too few for a comparison, and events dated only by month are left out except the Italian suspension of March 2023, dated 31 March, when the regulator's order was announced and access was blocked. The event day is the first trading day on or after the date of the incident. Of 44 such incidents, 38 have a full estimation window and at least three firms on each side: events 3 to 9, 11, 12, 14 to 16, 19, 22, 23, 26, 27, 31 to 34, 36 to 38, 40, 41, 44 to 47, 49 and 51 to 57.

Estimation and inference. For each incident and firm, a market model is estimated by least squares over trading days -250 to -30, with at least 120 returns, and abnormal returns are cumulated over each window. Within an incident, the mean over the clients of the affected provider is compared with the mean over firms whose products run only on other providers' models, and the test statistic is the mean of these differences across incidents divided by its standard error across incidents, which allows for any correlation among firms on the same day. The placebo draws, 1,000 times, as many dates as there are incidents from trading days between the first and last incident that lie more than five calendar days from any incident, assigns them the providers of the actual incidents in order, and records the mean difference

over days 0 and 1, with abnormal returns from a market model estimated over the whole sample for speed; the observed mean computed the same way is -1.31 points, against a placebo standard deviation of 0.46 points.

Relation to the earlier pooled test. A version of this paper compared the firms exposed to each incident with the market alone, using firms identified from all filing types that name OpenAI or Anthropic, and found an average abnormal return of -0.4% over days 0 and 1 with a t of -1.0. The clients of the affected provider here lose 0.9% ($t = -2.2$), of the same order. The comparison across providers adds power because the clients of different providers share most of what moves AI-exposed shares on any given day, and differencing removes it; it also shows that the clients of rival providers tend to gain on the day, which a comparison with the market alone counts as noise.

F Extensions

Limits per provider event. Between full cover and exclusion lies a cap L, as a share of a full-market event, on the losses paid for the failure of any one provider. The accumulation becomes $T_L = \sum_k \min\{\sigma_k, L\} F_k \le L\,N$, where $N = \sum_k F_k$ counts the failing providers, so that $\mathrm{ES}_\alpha(T_L) \le L\,\mathrm{ES}_\alpha(N)$ by the monotonicity and positive homogeneity of expected shortfall; the expected number of failing providers in the tail, $\mathrm{ES}_\alpha(N)$, lies between 0.9 and 1.8 at 99% for failure probabilities from 0.13 to 2%. A limit of the order of $r(v-\eta)/[\kappa\,\mathrm{ES}_\alpha(N)]$, between 0.14 and 0.27 of a full-market event at $r = 0.1$ and the lower threshold, therefore makes the accumulation insurable while leaving its tail with the insured firms, whose idiosyncratic harms remain covered in full. With a limit of a tenth of a full-market event, stand-alone capital for the end-2025 market at $p = 1\%$ falls from 0.40 to 0.10 and the index from 3.6 to 1.0, and the loss of the largest-provider scenario falls to 3.5% of premium, or 5.8% of the equity of a liability writer with equity of 60% of premium; 95% of the expected AI loss remains insured, while 75% of the loss from a failure of the largest provider is left with the insured firms. A limit of a fifth gives an index of 1.9, with half of that loss uninsured. With such limits, insureds can also keep idiosyncratic cover and buy part of the accumulation separately, valuing a marginal unit of it at $(v-\eta)\bar{p}$; in the sharing equilibrium of Proposition 2 the share covered is then the largest $c_T \le 1$ with $\kappa e_1(c_T/S) \le (v-\eta)\bar{p}$.

Nested concentration. Providers depend on a small number of cloud and chip suppliers. Suppose provider k runs on infrastructure $j(k)$ and that an infrastructure failure $G_j \sim \mathrm{Bernoulli}(p_j^{\mathrm{inf}})$, independent of everything else, disables all providers that run on it, so that provider k is down if $F_k = 1$ or $G_{j(k)} = 1$. Write $\sigma_j^{\mathrm{inf}} = \sum_{k:j(k)=j} \sigma_k$ for the share of adopters on infrastructure j and $\bar{p}_j = (\sigma_j^{\mathrm{inf}})^{-1} \sum_{k:j(k)=j} \sigma_k p_k$ for the average failure probability of its providers. Conditioning on the infrastructure states gives the systematic variance exactly:

$$\mathrm{Var}(T) = \sum_j (\sigma_j^{\mathrm{inf}})^2 (1-\bar{p}_j)^2 p_j^{\mathrm{inf}}(1-p_j^{\mathrm{inf}}) + \sum_k \sigma_k^2 (1-p_{j(k)}^{\mathrm{inf}})\, p_k(1-p_k).$$

The first term is a Herfindahl index of the infrastructure layer, weighted by failure variance. The second is the provider term of Lemma 1, slightly reduced because a provider's own

failure adds nothing when its infrastructure is already down. When infrastructure fails about as often as providers do, the two terms are of the same order; in the example of Appendix C, nesting roughly doubles the systematic variance. The tail bound of Lemma 2 carries over: expected shortfall is at least $\max_j \sigma_j^{\text{inf}} \min\{1, p_j^{\text{inf}}/(1-\alpha)\}$. Insurers already model the failure of a cloud provider as a scenario of its own (Lloyd's and AIR Worldwide 2018), and whenever providers share infrastructure, provider shares understate the relevant concentration. The same logic applies to other common causes: a regulatory measure such as the export-control directive of June 2026, which suspended two models at once and could as easily have reached several providers, or a vulnerability shared by models built on similar methods. The same decomposition applies when the common layer sits above the models rather than below them: an application vendor whose tool many deployers use, such as the screening tool in the litigation of Section 2.3, is a block of its own, whose failure reaches all its clients regardless of the model it runs on, and whose share enters the first term in place of σ_j^{inf}.

Ambiguity. The failure probabilities p_k are unknown. There is no loss history for liability-relevant failures of foundation models, and the failure modes change with every release. Suppose an insurer or a supervisor regards any p in an interval $[p^-, p^+]$ as plausible. If the interval straddles $1-\alpha$, value-at-risk capital for a concentrated market can be zero or close to the largest market share depending on which point in the interval is used, while expected shortfall moves continuously. Model uncertainty about p does not diversify across policies either; it is a common factor in its own right. An insurer that guards against it by evaluating the accumulation at the least favourable plausible value, p^+, holds more capital, but within a large and varied book the accumulation's integrated capital at p^+ still tends to zero as its relative size shrinks (Proposition 3). Treating the accumulation as a stand-alone scenario requires in addition an ambiguity loading charged peril by peril; under smooth ambiguity aversion, scenario-based pricing of this kind is the large-book limit of an insurer's problem when ambiguity does not pool away (Aslan 2026d).

Multi-homing. A firm that routes the same workflow to two providers and fails over when one is unavailable reduces its exposure to outages; a firm that uses different providers for different workflows does not, since each workflow still depends on one model. Failover does nothing against a defect that changes a model's outputs without making it unavailable, the case most relevant to liability. The relevant quantity is the share of workflows without an effective fallback, which is unobserved; the calibration treats q as a parameter for this reason.

Heterogeneous failure probabilities. Part (ii) of Lemma 2, part (i) of Lemma 3 and part (i) of Proposition 4 hold for any $p_1, \dots, p_K$. The Schur-convexity results need exchangeability. Without it, the systematic variance is $\sum_k \sigma_k^2 p_k(1-p_k)$, a concentration index weighted by each provider's failure variance, and concentration can lower capital if the larger providers are the safer ones (Appendix C).

Latent claims. The model lets all claims from a failure fall in the year in which it occurs. Many AI harms would be latent instead. A screening tool or an advice model used for years can generate claims that surface slowly, depend on how courts treat the first cases and reach

several lines of business at once. The closer analogues in insurance history are latent casualty accumulations such as asbestos, in which one common cause produced claims across many policy years and lines and took decades to resolve (Carroll et al. 2005). Latency mainly affects when the loss is recognised. For insurers it moves the accumulation from catastrophe capital into reserve risk, whose standard-formula factors are equally silent about the technology insureds share, and it can stack the same exposure across several underwriting years; for banks it spreads defaults over several years (Section 5).

Pricing dependence. If firms choose their providers and insurers cannot observe which provider an insured depends on, premiums are pooled and firms do not face the correlation that their choice creates. At end-2025 shares and $p = 1\%$, the marginal contribution of each provider's clients to the stand-alone expected shortfall of the accumulation, net of expected loss, is 0.99 per unit of share for the largest provider and at most 0.02 for the others. Pricing by contribution would therefore charge clients of the largest provider capital costs of 5.9% of a year's premium at $z = 1$ and a 6% cost of capital, against 2.4% for every adopter under pooling, and almost nothing to clients of the other providers. A charge of this kind would give firms an incentive to spread their dependence across providers that pooled premiums do not give, but it requires the data on dependence discussed in Section 7.

Reinsurance and capital markets. Reinsurers cannot diversify T away within the AI books they assume, since every cedent's exposure loads on it, but they hold it alongside unrelated perils, which is what makes its incremental capital small (Appendix C.3). Capital markets can diversify it only to the extent that investors do not already hold the same factor through AI-related equity and debt, as insurers themselves do (see the asset side below). Securitisation runs into a problem of trigger design: an indemnity trigger requires the claims to have been settled, and a parametric trigger on a provider event leaves cedents with basis risk, a problem I study for systemic cyber risk in Aslan (2026c).

The asset side. Section 5.3 computes the combined loss for banks that hold the provider factor on both sides of the balance sheet. The same holdings are spread more widely: US life insurers held USD 85.2 billion of investments linked to the AI ecosystem in mid-2026, according to Moody's as reported in the trade press (Reinsurance News 2026). For insurers, such holdings are the case of correlated sharing in Appendix C.3: if the asset loss is non-decreasing in T, it is comonotonic with the liability losses of AI-exposed books, and expected shortfall adds the two without diversification (Dhaene et al. 2006).

G Upstream liability under the revised Product Liability Directive

This appendix models how the revised Product Liability Directive shifts part of the excluded loss to providers, gives the legal background to the share of each claim that reaches them, calibrates the model and proves its result.

G.1 Model

The Directive makes the developer of a defective model liable without fault for harm to natural persons, jointly and severally with the manufacturer of the product into which the model is integrated. Write μ_D for the part of the expected loss per harm that arises from harm the Directive covers. Strict liability, presumptions of defectiveness and causation, and disclosure duties raise the expected claims per harm from μ_D to $\beta\mu_D$, with $\beta \geq 1$, and joint and several liability and recourse allocate a share $\lambda \in [0, 1]$ of each claim to the provider whose model failed, a share that the caps and indemnities in providers' terms (Section 2.3) leave uncertain. In the failure state, provider k is liable for

$$Q_k = \lambda\beta\mu_D q\, a\, \sigma_k \tag{13}$$

and can pay at most its loss-absorbing capacity in that state, A_k, its liquid equity impaired by the failure plus the limit of its own liability insurance, so that its recovery ratio is $\tau_k = \min\{1, A_k/Q_k\}$. What it cannot pay returns to downstream manufacturers and their insurers, or stays with the injured person where the deployer supplies a service.

Proposition 4 (Upstream liability)**.** *(i) For given capacity, the recovery ratio τ_k is non-increasing in the provider's market share σ_k.*

(ii) If $A_k = \bar{A}$ and $p_k = p$ for all k, the expected unrecovered liability $\sum_k p_k\,(Q_k - \bar{A})^+$ is Schur-convex in σ: it does not fall when concentration rises.

(iii) If $A_k = \chi\sigma_k$, the expected unrecovered liability is $(\lambda\beta\mu_D qa - \chi)^+\,\bar{p}$, which with equal failure probabilities does not depend on concentration and is zero if and only if $\chi \geq \lambda\beta\mu_D qa$.

(iv) Downstream manufacturers and their insurers bear $\beta\mu_D(1 - \lambda\tau_k)$ per harm after the reform, against μ_D before it. If $\beta(1-\lambda) \geq 1$, the reform lowers this burden for no client. If $\beta > 1 > \beta(1-\lambda)$ and $A_k = \bar{A}$, it lowers the burden for clients of providers with share below

$$\sigma^* = \frac{\bar{A}}{(\beta - 1)\,\mu_D q\, a} \tag{14}$$

and raises it for clients of larger providers.

Liability shifts risk only as far as the liable party's resources (Shavell 1986; Pitchford 1995; Boyer and Laffont 1997). With concentrated providers, the claims against each provider are proportional to its market share and fall due in the state in which its own value has fallen. With capacity of 5% of the claims that a failure of all providers would generate and $\lambda = 0.2$, the largest provider, at a share of 40%, could pay 63% of what it owes. If the reform raises claims by half ($\beta = 1.5$), its clients bear 31% more per harm than before, and even clients of providers small enough to pay in full bear 20% more, because the reform enlarges claims by more than it moves upstream. Part (ii) applies when capacity is set by the market's per-risk insurance capacity and not by the provider's size; when capacity grows with share, the concentration effect disappears, and part (iii) points to a remedy: a financial responsibility requirement of $\lambda\beta\mu_D qa$ per unit of market share makes recoveries complete and independent of concentration.

The allocation share. The direction in which the Directive moves λ is not settled. An injured person can claim the whole loss from a solvent provider, and liability towards the injured person cannot be limited by contract (Art. 15). Recourse between businesses, however, is left to national law, where the caps and indemnities in providers' terms apply and the indemnities run from the customer to the provider (Appendix A.4); a provider can also attribute a defect to the deployer's design or instructions (Art. 11(1)(f)), and the development-risk defence (Art. 11(1)(e)) may cover failure modes that were not discoverable when the model was released. What a provider cannot pay returns, through joint and several liability, to downstream firms that are themselves manufacturers of a product and to their insurers; where the deployer supplies a service, the unpaid share stays with the injured person unless a national compensation scheme pays (Art. 8(5)). The comparison in Proposition 4 takes the pre-reform allocation to be entirely downstream; national rules on contractual and fault-based liability, which the Directive leaves in place (Art. 2(4)(b)), already shift some of the loss upstream.

G.2 Calibration

I set $\lambda = 0.2$ because a cap of a year's fees allows recourse for at most about a quarter of the claims at $z = 1$: enterprises spent USD 12.5 billion on model interfaces in 2025 (Tully et al. 2025b), about a quarter of the US liability premium written on adopters at the central adoption rate, which sets the scale of claims in a full-market event at $z = 1$. The comparison is one of orders of magnitude, since the Directive's claims are European and the premium figure American. The upper case $\lambda = 0.7$ corresponds to courts setting the caps aside as unfair terms when the provider's own model was defective, so that recourse is limited only by the deployer's share of fault in integrating and supervising the model. Claims expansion factors are $\beta = 1.25$ and 1.5, and provider capacity $\hat{A} = \bar{A}/(\beta\mu_D qa)$, a share of the claims that a failure of all providers would generate, lies between 2 and 10%. With $\lambda = 0.2$ and $\beta = 1.25$, $\beta(1-\lambda) = 1$ exactly, so the central case sits on the boundary of part (iv), and $\beta = 1.5$ is the more informative case. With capacity of 5% of full-market claims, a provider can pay in full up to a market share of 25%, and at the largest provider's 40% its recovery ratio in the failure state is 63%. With $\beta = 1.25$, the burden per harm is unchanged for clients of providers with shares up to 25% and rises by 2 and 9% for clients of the second and the largest provider; with $\beta = 1.5$, it rises by 20% for clients of every provider up to that share and by 22 and 31% for the two largest. With capacity of 2%, the largest provider's clients bear 19 or 43% more. Under the upper bound $\lambda = 0.7$, clients of small providers gain and the burden rises only above a share of 15 to 25% (Figure 10). Two of the three largest providers at end-2025 are stand-alone developers, whose capacity in a failure consists of their own liquid resources and their liability insurance, reported at up to USD 300 million for one of them (Section 2.4); the third belongs to a technology group whose balance sheet is of a different order. These shifts reach downstream firms that manufacture products; where the deployer supplies a service, the unpaid share stays with the injured person.

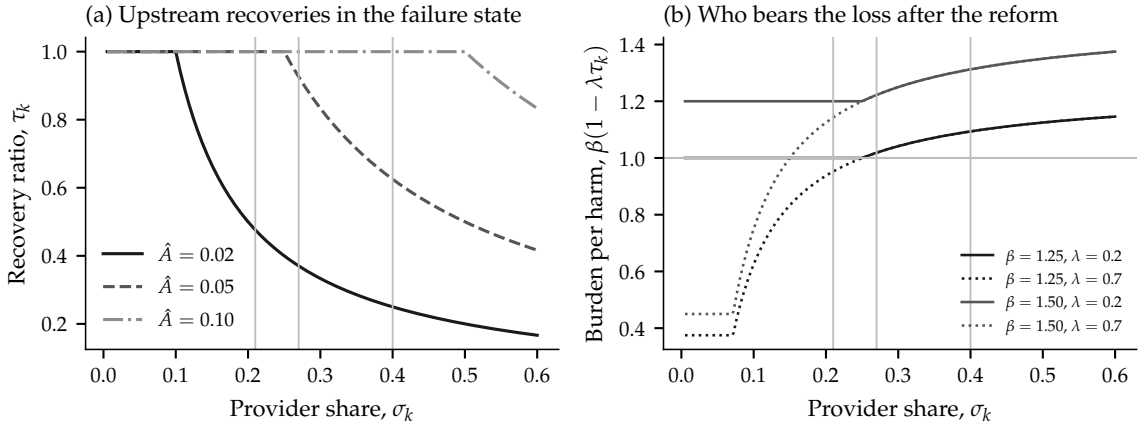


Figure 10. Recoveries from providers and the downstream burden under strict product liability
Notes: Panel (a): recovery ratio $\tau_k = \min\{1, \hat{A}/(\lambda\sigma_k)\}$ for capacity $\hat{A}$ expressed as a share of the claims that a failure of all providers would generate, $\lambda = 0.2$. Panel (b): downstream burden per harm after the reform relative to before, $\beta(1 - \lambda\tau_k)$, with $\hat{A} = 0.05$, for $\lambda = 0.2$ (solid) and $\lambda = 0.7$ (dotted). Vertical lines mark the end-2025 shares of the three largest providers.

G.3 Policy

European law has precedents for pairing liability with financial security: the Environmental Liability Directive asks Member States to encourage financial security instruments for operators (European Parliament and Council 2004), the European Parliament proposed compulsory liability insurance for operators of high-risk AI systems in 2020 (European Parliament 2020), and the Directive itself allows compensation schemes for victims of insolvent operators (Art. 8(5)). Because the Directive harmonises liability itself, a financial responsibility requirement of $\lambda\beta\mu_D qa$ per unit of market share would fit better in the AI Act, alongside the duties it already places on providers of general-purpose models with systemic risk. Whether national rules on unfair standard terms limit the caps and indemnities in providers' terms when the provider's own model was defective is untested; a common rule on recourse for defects in general-purpose models would settle it.

Proof of Proposition 4. (i) The liability Q_k in (13) is increasing in σ_k, so for given capacity A_k the recovery ratio $\tau_k = \min\{1, A_k/Q_k\}$ is non-increasing.

(ii) With $p_k = p$ and $A_k = \bar{A}$, the expected unrecovered liability is $p\sum_k f(\sigma_k)$, where the function $f(x) = (\lambda\beta\mu_D qa\, x - \bar{A})^+$ is convex. A sum of the same convex function of each coordinate is symmetric and convex, hence Schur-convex.

(iii) With $A_k = \chi\sigma_k$, $(Q_k - A_k)^+ = \sigma_k(\lambda\beta\mu_D qa - \chi)^+$, and $\sum_k p_k\sigma_k = \bar{p}$, which equals p for every σ when $p_k = p$.

(iv) If $\beta(1-\lambda) \geq 1$, then $\beta\mu_D(1-\lambda\tau_k) \geq \beta\mu_D(1-\lambda) \geq \mu_D$ for every k and every capacity. Otherwise, let $\hat{\sigma} = \bar{A}/(\lambda\beta\mu_D qa)$. For $\sigma_k \leq \hat{\sigma}$, $\tau_k = 1$ and the downstream burden is $\beta\mu_D(1-\lambda) < \mu_D$. For $\sigma_k > \hat{\sigma}$, $\tau_k = \bar{A}/(\lambda\beta\mu_D qa\sigma_k)$ and the burden is $\beta\mu_D - \bar{A}/(qa\sigma_k)$, which is increasing in σ_k and exceeds μ_D if and only if $\sigma_k > \sigma^*$ when $\beta > 1$. Finally, $\sigma^* > \hat{\sigma}$ if and only if $\lambda\beta > \beta - 1$, that is, $\beta(1-\lambda) < 1$. □

H Supplementary results

This appendix collects results summarised in the main text: the agents of credit facilities of Section 2.1, the reverse stress test and the system-wide loss of Section 5, the financing of AI infrastructure in filings, the syndicated lenders and the sourcing changes of Section 6, and the predictions of Section 7.

H.1 Agents of credit facilities of firms that name a provider

Where firms do name one, their lenders can be traced. Reading the filings that name OpenAI or Anthropic as a dependency between January 2025 and October 2026 identifies 93 firms, 80 of them naming OpenAI, 29 Anthropic and 16 both, of which 71 state that the model is used in their own products or operations rather than mentioned in passing. Seventeen are foreign private issuers filing on Forms 20-F or 6-K, among them SAP SE, Genmab, Immatics, Sapiens and the UK holding companies of a euro-area banking group, so the disclosure is not confined to US firms even though the search is of US filings. Of the 93, 32 disclose a credit facility and 31 name its administrative agent. Those agents are concentrated: one bank is agent to 11 of the 31, the three largest to two-thirds, and their Herfindahl index is 0.19, close to the lower end of the 0.21 to 0.37 range for provider shares themselves (Table 12 in Appendix C); two of the facilities are agented by private credit funds rather than by banks. A concentration of lenders therefore sits above the concentration of providers. The chain then stops where it would matter: an administrative agent arranges and administers a syndicate without disclosing how much of the loan it retains, the other members are rarely named, and no filing records in one place which model a borrower depends on and which institution holds the resulting credit risk. These are US filers whose agents are mostly US banks, a different population from the European banks of Section 5.2, so the two are parallel illustrations of the same gap rather than a single chain. Section 6 follows a larger set of firms, identified from annual reports, to every member of their lending syndicates and the share each holds.

H.2 A reverse stress test

Because the loss is the product of the ratio of corporate loans to capital, the share x of those loans that go to clients of the failing provider, the loss given default and $q\Psi_0$, the calculation can be inverted (Figure 11a). A bank loses 1% of its CET1 capital in the failure year once $x\,q\Psi_0$ exceeds 0.69% if it has the median ratio of corporate loans to capital, 0.40% at the ninetieth percentile of the ratio and 0.24% at the highest ratio, which at the central $q\Psi_0$ of 1.5% means 46, 27 and 16% of corporate loans. At the average adoption rate, with 35% of corporate loans going to AI adopters, the bank with the median ratio would not reach the threshold even if all its adopting borrowers relied on the failing provider; it would lose 1% only if $q\Psi_0$ reached 4.9% with borrowers that mirror the market or 2.0% with adopting borrowers all on the failing provider. For the bank at the ninetieth percentile, dependence of about three-quarters of its adopting borrowers on the failing provider would suffice, and for the bank with the highest ratio a little under half, close to the provider's market share of 40%. A tenth of the headroom above requirements is used once $x\,q\Psi_0$ exceeds a threshold whose median across banks is 2.4%, with one bank in ten below 0.94% and the lowest at 0.42%. Supervisors can compute

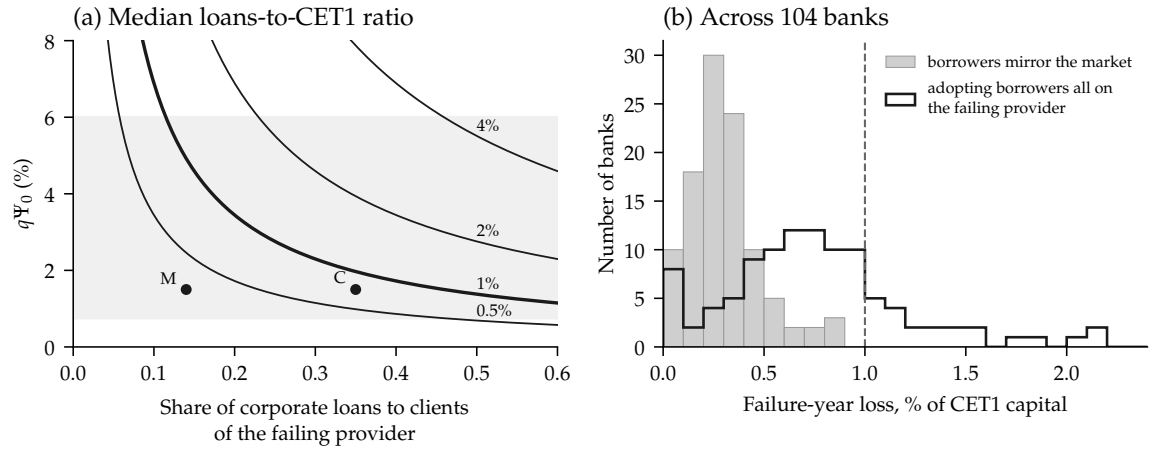


Figure 11. Failure-year loss as a share of CET1 capital

Notes: Panel (a): lines of equal loss in a year in which the largest provider fails, loss/CET1 $=$ (NFC loans/CET1) $\times$ $x \times$ LGD $\times$ $q\Psi_0$, for the bank with the median ratio of loans to non-financial corporations to CET1 capital, 3.6, at losses of 0.5, 1, 2 and 4% of CET1 capital, where x is the share of those loans that go to clients of the failing provider and LGD $= 40\%$; at the average adoption rate, M marks the central scenario for borrowers that mirror the market ($x = 0.35 \times 0.40$) and C for adopting borrowers all on the failing provider ($x = 0.35$). The shaded band is the range of $q\Psi_0$ on the grid, 0.75 to 6%. Panel (b): failure-year loss of the 104 banks of Table 20 (European Banking Authority 2025) at $q = 10\%$ and $\Psi_0 = 15\%$, with each bank's adoption rate scaled by the sector composition of its loans; the dashed line marks a loss of 1% of CET1 capital.

these thresholds from the loan and capital data they already collect, but they cannot place a bank against them without knowing x.

H.3 The banking system

Euro-area significant institutions held EUR 6.2 trillion of loans to non-financial corporations and EUR 1,560 billion of CET1 capital in mid-2026 (European Central Bank 2026b), a population that differs from the EBA sample. In the year the largest provider fails they would lose EUR 5 billion, or 0.33% of their CET1 capital, and 0.2 to 1.3% over the grid. If all adopting borrowers relied on the failing provider, the same association between capital and lending, applied to each of the 104 banks of Section 5.2, puts loan growth in the year after the failure 0.03 percentage points lower at the median bank and at most 0.08 points lower at any bank. No bank would exhaust its buffer above requirements. If the external cost Γ of Proposition 1 rises steeply only once losses exceed such buffers, as Section 3.4 suggests, the wedge W is small at these values; I do not quantify Γ, so this is a reading of the magnitudes rather than an estimate. For the system as a whole, the case for supervisory attention rests on lenders whose borrowers' dependence is concentrated. The conclusion assumes independence; if provider failures coincide with downturns, so that the macroeconomic factor is one standard deviation worse in failure years, exclusion raises the 99.9% quantile of the corporate portfolio by 5 to 39 basis points against 3 to 26 of expected loss, and economic capital rises by up to 13 basis points, or 2.4% of the formula's capital at the original default probability.

H.4 Financing of AI infrastructure in filings

The firms that build this infrastructure file with the SEC, so their borrowing can be read directly even though banks' holdings of it cannot. Ten US-listed AI-infrastructure firms

disclose 31 financings since 2025: 22 loan facilities and nine note issues, USD 49.9 billion in all, USD 33.0 billion of it in loans and USD 16.9 billion in notes, with one bridge facility whose size is not disclosed (Table 13 in Appendix C). Sixteen of the loan facilities have a commercial bank as administrative agent, and two banks arrange eleven of the sixteen. JPMorgan Chase is agent to CoreWeave's revolving and letter-of-credit facilities and one of its delayed-draw facilities, to Core Scientific's revolver and letter-of-credit facility and to Hut 8's revolver; Morgan Stanley is agent to another CoreWeave delayed-draw facility, to Cipher Mining's revolver, to TeraWulf's revolver and bridge and to Core Scientific's bridge. Both appear again on the liability channel, where JPMorgan Chase is agent to 11 of the 31 facilities of firms that name a provider and Morgan Stanley to 4, behind Bank of America with 5; across the two channels the two of them are agents to 26 of the 45 facilities that have a bank as agent. The comonotonicity of equation (11) is thus not only a property of the loss distribution: the credit on both sides is arranged by the same houses.

In their announcements, five of these financings state the contract that services them. IREN's USD 1.5 billion delayed-draw facility and USD 2.1 billion of notes are raised against an agreement with Microsoft to provide dedicated GPU services at data centres in Childress, Texas (IREN Limited 2026); Cipher Mining's notes rest on a lease whose tenant holds a recognition agreement with Google; Galaxy Digital's notes name CoreWeave as the tenant; and Applied Digital's notes reference an assignment of a CoreWeave lease. The announcements of CoreWeave's own five delayed-draw facilities, USD 19.8 billion, say only that they finance capital expenditure required to perform customer contracts, and name no customer. The credit agreements of two of them, filed as exhibits, do: their covenants turn on payments by OpenAI and Microsoft and on OpenAI's credit rating (Section 6.5). The information exists, but only in the exhibits and only for the borrowers that file them, which is the same gap in disclosure that Section 2.1 finds among deployers, one layer further up.

The range for ι is nonetheless wide, because what can be read from filings is the borrower's side and not the lender's. An administrative agent does not disclose the share of a facility it retains, the other syndicate members are not named, and the notes are held by investors the filings do not identify, so USD 49.9 billion of identified financing cannot be allocated to any bank's balance sheet from filings alone; for the part raised as syndicated loans, commercial loan data record the syndicate and most of the shares (Section 6.2). No supervisor publishes lending to AI infrastructure by bank, and the exposure falls into different NACE sections depending on how a data centre is incorporated and financed, so it cannot be recovered from the sectoral breakdown used in Section 5.2 either. The quantity supervisors would need has the same shape as the one they lack on the liability side, the share of exposures tied to each provider, and the ad hoc question proposed in Section 7 would answer both at once. Neither exposure is aggregated by provider in published supervisory data.

H.5 Syndicated lenders and sourcing changes

Table 27 lists the largest holders of syndicated facilities to the firms that disclose dependence (Section 6.2), and Table 28 the providers named in successive annual reports (Section 6.4).

Table 27. Syndicated lending to firms that disclose dependence, 30 June 2026

Lender	Country	Rank by book	Exposure (USD bn)	Share of book (%)	Infrastructure (% of book)
Bank of America	US	2	3.43	0.53	0.00
JPMorgan Chase	US	1	3.27	0.46	0.23
Royal Bank of Canada	Canada	17	2.41	0.85	0.00
BNP Paribas	France	8	2.31	0.47	0.00
Wells Fargo	US	5	2.23	0.43	0.29
Citigroup	US	4	2.05	0.38	0.22
Morgan Stanley	US	13	2.01	0.60	0.50
Truist Financial	US	21	1.94	0.93	0.00
Mitsubishi UFJ	Japan	3	1.70	0.30	0.26
PNC	US	22	1.70	0.83	0.72
Ares Capital	US	59	1.50	2.67	0.00
U.S. Bancorp	US	20	1.49	0.69	0.14
All 86 lenders with exposure			50.0	0.28	0.15

Notes: The twelve largest holders of syndicated facilities to the 22 borrowers that disclose in annual reports that their products run on a provider's model or that they use one internally, among tranches active on 30 June 2026 (LSEG 2026). Exposure: commitment times the lender's share, with missing shares imputed as described in the text. Rank by book: rank among about 4,000 lenders by total syndicated exposure on the same date. Share of book: exposure to these borrowers as a share of the lender's total syndicated exposure. Infrastructure: exposure to CoreWeave, its financing vehicles, Galaxy Digital's Helios subsidiary and TeraWulf as a share of the same total.

Table 28. Providers named in successive annual reports of firms that disclose dependence

	Annual	Naming a	Naming several		Consecutive	Provider	Provider
Filed in	reports	provider	Reports	%	pairs	added	dropped
2023	12	9	0	0	0	0	0
2024	48	40	3	8	5	0	0
2025	71	62	9	15	35	4	0
2026	87	80	25	31	48	6	2
2023–2026	218	191	37	19	88	10	2

Notes: Annual reports, excluding amendments, of the 114 firms that state that their products run on a provider's models or that they use one internally. A report names a provider if a sentence names it, or one of its models, outside a description of competitors or of the market. Consecutive pairs: pairs of successive reports of the same firm in which the earlier one names a provider; the year is that of the later report. Additions and drops are checked by hand against the passages and listed in the replication files.

H.6 Predictions

The explanations in Table 23 (Appendix D) can be told apart once dependence is observed. The accumulation predicts that the adoption and tightening of AI exclusions rise with the share of a carrier's insureds whose workflows depend on the largest provider, scaled by the carrier's capital, that they are stronger at carriers that set appetites scenario by scenario and lack reinsurance for AI (Propositions 2 and 3), and that affirmative cover is backed by diversified reinsurers and, once claims can be attributed to a provider, comes with aggregate limits per provider event (Appendix F); contract certainty predicts none of these differences across carriers. Rate and form filings with US state regulators can be matched to carriers' books by line. For banks, the provider factor should appear in losses in the year of a failure and be largest for lenders concentrated on one provider, while risk weights respond only as

default estimates are updated; in the euro area, loan-level data in AnaCredit could be matched to firms' AI use within the secure environments in which enterprise survey microdata can be accessed. A defect in a model's outputs would test the mechanism directly, through the claims, defaults and abnormal returns it produces among the clients of the failing provider; the incidents on record already move those clients' shares relative to the clients of other providers (Section 6.3), while their loan spreads at origination and their bond prices and yields show no premium (Section 6.4); the same comparisons, applied to default rates and to banks' own exposures, are the tests that wider dependence data would allow for lenders. The covenants in the credit agreements of Section 6.5, which turn on a provider's payments and rating, show that the information is already collected loan by loan; what is missing is its aggregation by provider.