Data as Liability

Fatih Kansoy*
University of Oxford

Yuhao Huo**
University of Cambridge

This version: March 25, 2026

Abstract

Artificial intelligence makes data more productive, but it also makes data more costly to govern. This paper asks where that governance cost shows up in firms' own risk disclosures. Using roughly 84000 firm-years of SEC annual filings for US listed firms from 1994 to 2023, the paper builds layered text and LLM measures to separate AI invention from AI adoption and relates both to disclosed attention to data breach risk. AI adoption is associated with roughly 5 per cent higher breach-risk attention relative to the sample mean; AI invention is economically negligible once both margins enter the same specification. The wedge survives an explicit non-AI digitisation placebo built from the same filings. Among adopters, breach-risk attention is highest where deployment is customer-facing. Firms that explicitly connect AI to breach vulnerability describe it as expanding exposure in 101 of 103 directional statements. Supplementary evidence from staggered state Data Breach Notification laws is directionally consistent with the disclosure results.

Keywords: Artificial intelligence, data breach risk, corporate disclosure, Data Breach Notification laws, text-as-data, staggered difference-in-differences

JEL Classification: O33, G38, D83, L86, K22

*Corresponding Author. Department of Economics, Oxford. fatih.kansoy@economics.ox.ac.uk
**Department of Economics, Cambridge. Contact: yh620@cam.ac.uk

1. Introduction

Artificial intelligence is usually discussed as a technology that turns data into value. For firms, however, the same process can also turn data into a liability. The systems that make prediction cheaper and personalisation easier require firms to collect, move, store, and secure larger volumes of sensitive information across more products, vendors, and operational links. This paper studies disclosed breach-risk attention and asks whether that salience is more concentrated in AI deployment, especially customer-facing deployment, than in AI invention.

The central distinction is between building AI capability and operating AI through a live data process. A firm can invest in models, algorithms, and technical capability without necessarily processing large volumes of personal information in day-to-day operations. Running AI in live products is different. Once AI is embedded in recommendation systems, customer service, underwriting, fraud detection, or targeted marketing, the firm is not only developing a technology; it is running a recurring data process. This raises three related questions. Is AI adoption, as distinct from AI invention, more strongly associated with firms' disclosed attention to data breach risk? Is that relationship strongest in customer-facing deployment? And does regulation that raises the expected cost of a breach align more closely with deployment than with research?

These questions connect three literatures that have mostly progressed in parallel, namely the economics of AI diffusion (Babina et al. 2024, Agrawal, Gans, and Goldfarb 2019), the corporate finance of cybersecurity risk (Florackis et al. 2023, Kamiya et al. 2021, Jamilov, Rey, and Tahoun 2025), and the political economy of data regulation (Miller and Tucker 2009, Goldfarb and Tucker 2012, Frey, Presidente, and Andres 2024). The outcome is *disclosed* breach-risk attention, not realised breach incidence. The distinction is substantive, not incidental. Risk disclosures reveal managerial salience and materiality in the firm's stated business model, and prior work shows that this type of disclosure contains economically meaningful information (Verrecchia 1983, Dye 1985, Campbell et al. 2014, Florackis et al. 2023). The analysis asks when AI diffusion makes breach risk salient enough to be disclosed; it does not ask whether every disclosure maps one-for-one into a realised incident.

To answer those questions, the analysis uses SEC 10-K filings for US listed firms from 1994 to 2023. The resulting panel covers roughly 86,000 firm-years and more than 10,000

firms. The Business section of the 10-K (Item 1) provides the text-based measures of AI invention and AI adoption. The Risk Factors section (Item 1A) provides the filing-based measure of breach-risk attention. The mechanism analysis combines filing-level and paragraph-level AI-use classifications with extracted risk statements from firms' own disclosures. The regulatory analysis uses the staggered enactment of state Data Breach Notification laws as supplementary evidence on whether regulation affects research and deployment differently.

The main disclosure evidence yields a clear asymmetry. Firms with more AI adoption devote more of their risk disclosures to breach-related concerns, whereas firms whose AI language is concentrated in invention do not in the lead filing-level horse race when both margins enter together. The wedge survives an explicit control for non-AI digitisation language in the same Business section text. It is also strongest for customer-facing deployment, where recurring personal-data handling is hardest to separate from the technology itself.

Firm and year fixed effects, filing-length controls, and within-between decompositions sharpen the comparison, but they do not rule out selection into AI adoption, a limitation the conclusion addresses directly.

Firms' own disclosures reinforce that reading. When they explicitly connect AI to cyber or breach vulnerability, they almost always describe AI as expanding exposure, not reducing it. The DBN evidence is supplementary, but its estimated signs line up with the disclosure results and echo the main disclosure evidence.

The evidence therefore supports an economic contribution rather than only a measurement one. In deployment, data is both a productive input and a channel of liability. The invention-adoption wedge, the customer-facing mechanism, and the supplementary DBN evidence are consistent with that interpretation. The managerial implication is that deployment decisions bundle productivity gains with governance costs. For policymakers, the evidence is consistent with data-centred regulation aligning more closely with deployment margins than with research margins.

The rest of the paper is organised as follows. Section 2 sets out the conceptual framework. Section 3 reviews the related literature. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports the results. Section 7 concludes.

2. Conceptual Framework

At the centre of the framework is the distinction between AI invention and AI adoption. Some AI activity expands technical capability; other AI activity embeds that capability in a live data process. The latter should matter more for breach-risk salience because it requires the firm to handle data repeatedly in products, services, and operational workflows. Because the outcome is disclosed breach-risk attention, the framework concerns revealed materiality and managerial salience rather than realised incidents alone. Three comparative predictions follow: one about the difference between invention and adoption, one about heterogeneity within adoption, and one about how regulation that raises the expected cost of a failed data process should bear on those two margins.

The first and most important distinction is between *AI invention* and *AI adoption*. Invention includes activities such as model development, experimentation, and patenting. Operating AI is different. Once AI is deployed in recommendation systems, customer service, underwriting, personalisation, identity verification, or marketing, the firm must repeatedly gather, store, combine, and transmit the information that makes those systems useful. The exposure created by deployment persists for as long as the application remains in operation. The key contrast is between firms that describe AI as a research capability and firms that describe it as an operating technology.

Prediction 1. AI adoption is positively associated with disclosed attention to data breach risk; in the lead filing-level horse race, AI invention is not systematically associated with it when both margins enter together.

The second distinction is within adoption itself. Not every deployed AI system handles the same type of data or creates the same exposure. Customer-facing AI should be the most exposed form because it most directly processes customer identities, transactions, account histories, communication records, browsing behaviour, location data, or other individual-level information. Internal-use systems, third-party tools, and pure development can also matter for cyber risk, but they are less tightly tied to the recurring handling of customer data that activates legal, operational, and reputational consequences after a breach. The framework predicts a ranking, not an all-or-nothing distinction.

Prediction 2. Among firms with substantive AI activity, customer-facing AI is more strongly associated with disclosed attention to data breach risk than internal-use AI, third-party

adoption, or pure development.

A supplementary step introduces regulation. DBN laws raise the expected cost of a failed data process through notification expenses, legal exposure, reputational damage, and remediation burdens. They do not ban AI, but they make data-intensive deployment more costly to operate. The comparative prediction is straightforward. Regulation that raises the expected cost of a breach should bear more heavily on deployment than on the research margin.

Supplementary prediction. Regulation that raises the expected cost of a breach should bear more heavily on AI adoption than on AI invention.

The framework does not predict exact magnitudes, and it does not require every sector to look alike. The claim is comparative. Adoption should matter more than invention for disclosed breach-risk attention, customer-facing adoption should matter most within adoption, and regulation that raises the expected cost of a breach should weigh more heavily on deployment margins that rely on continuing data flows.

The disclosure analysis tests Predictions 1 and 2 by relating AI invention, AI adoption, and customer-facing deployment to breach-risk attention in firms' own risk disclosures. The DBN analysis then asks whether a legal shock that raises the cost of breach consequences feeds back onto adoption more than invention. The next step is to ask how far adjacent literatures already go and where the gap remains.

3. Related Literature

The gap spans three bodies of work. Research on AI diffusion has established why firms invest in AI and where returns show up, including growth, product innovation, productivity gains, and changes in organisational design (Babina et al. 2024, Brynjolfsson, Li, and Raymond 2023, Acemoglu et al. 2022, Felten, Raj, and Seamans 2021). That focus is natural. AI creates value by improving prediction, automating routine decisions, and helping firms scale products and services that depend on richer data. Yet the same feature that generates much of that value also creates exposure. Recommendation systems, automated underwriting, personalisation tools, and AI-enabled customer service are commercially useful precisely because they sit close to customer data, customer records, and repeated

data transfer. Less explored is the liability channel that follows when deployment requires firms to govern those recurring data flows.

A second omission matters as well. Much of the empirical AI literature works with broad measures of AI intensity that can bundle patents, hiring, software development, and product adoption into a single object. That aggregation is sensible for questions about a firm's place in the AI economy. It is less well suited to a question about data liability, because developing algorithms or filing patents need not require the live handling of personal customer information, whereas deployed systems often do. The distinction between invention and adoption therefore does more than rename two forms of AI activity. It separates a research margin from a data-processing margin. Recent work by McLemore and Mihov (2026) and Rishabh, Mihet, and Jang-Jaccard (2025) already points towards that difference, showing that AI can increase operational or cyber exposure when governance and deployment do not move together. An open question is whether that distinction appears in firms' own risk disclosures and whether regulation bears on the two margins differently.

The cyber-risk and disclosure literature supplies the outcome side of the paper. Data breaches are economically important for firms, investors, and customers (Kamiya et al. 2021, Akey et al. 2024, Foerderer and Schuetz 2022). A related disclosure literature shows that corporate risk disclosures are useful, not empty boilerplate. In voluntary-disclosure models, silence itself carries information, which gives firms incentives to discuss even unfavourable exposures (Verrecchia 1983, Dye 1985). Empirically, firms facing greater risk do disclose more risk factors, and those disclosures matter for pricing and market reactions (Campbell et al. 2014, Florackis et al. 2023). Topic models and related text methods have therefore become a credible way to measure risk salience from 10-K Risk Factors disclosures (Item 1A) (Bao and Datta 2014, Hassan et al. 2019, Sautner et al. 2023, Bybee et al. 2024). The literature is thinner on determinants. Disclosure-based cyber measures are often used as predictors of attacks or as pricing variables. Much less attention has gone to the technology choices that may cause some firms to discuss breach risk more than others.

A third strand, on privacy and data regulation, shows that data-dependent technologies do not diffuse in a legal vacuum. The canonical result is Miller and Tucker (2009), who show that privacy regulation slowed electronic medical-record adoption. Later work on digital advertising, data sharing, venture funding, and AI innovation reaches a related conclusion that when the cost of handling sensitive data rises, the adoption path and even

the composition of innovation can change (Goldfarb and Tucker 2012, Frey, Presidente, and Andres 2024). Less settled is whether research and deployment move together under those shocks. Some innovation margins can adjust by changing methods, inputs, or project direction. A deployed customer-facing system tied to ongoing personal-data handling has less room to evade the higher expected cost of a failure. DBN laws are useful in this setting because they do not prohibit AI use directly. They raise the expected cost of a failed data process through notification duties, legal exposure, remediation burdens, and reputational damage.

These literatures leave a specific opening. The gains from AI diffusion, the costs of cyber failure, and the effects of data regulation on data-dependent technology are all increasingly well understood. Much less is known about how those forces meet inside the firm. The central claim here is that disclosed breach-risk attention should be more tightly associated with deployment, especially customer-facing deployment, than with research. That claim calls for data that can separate invention from adoption, identify customer-facing use, and recover breach-risk salience from firms' own disclosures. It also gives the empirical strategy a more focused task than simply asking whether "AI" and cyber risk move together.

4. Data

Answering that question requires a filing-level panel from SEC EDGAR, augmented with supplementary measurement and validation layers. The filing corpus is the backbone of the design because the paper's central question is disclosure-based. The Business section captures what the firm says it does, and the Risk Factors section captures what it says its material risks are. Linking the two sides of the filing lets the paper connect AI activity to breach-related salience without relying on one narrative in one section. Table 1 is a simple map of the data structure. It shows which source produces each main layer and what that layer contributes to the analysis. The section then explains how those layers are constructed and why the filing-based measures remain the core design.

4.1 Filing Corpus and Sample Construction

Empirically, the starting point is a merged panel of 10-K filings for US listed firms from 1994 to 2023. After section extraction and identifier matching to CRSP-Compustat through CIK and related firm identifiers, the full filing panel contains 86,410 firm-years and 10,661 firms. Non-empty Business section text is available for 84,012 firm-years and 10,456 firms. Risk Factors text is available for 51,771 firm-years and 6,981 firms. Because Risk Factors disclosure becomes broadly available only after the SEC's 2005 mandate, the effective disclosure-analysis window begins in 2005 even though a smaller set of voluntary disclosures appears earlier.

TABLE 1. Core data overview

Layer	Main role	Level	Coverage
Panel A. Core design layers			
Business section text	Main AI invention and AI adoption measures	Firm-year	1994–2023; 84,012 firm-years / 10,456 firms
Risk Factors text	Main breach-risk attention outcome	Firm-year	1997–2023; main analysis 2005–2023; 51,771 firm-years / 6,981 firms
DBN law panel	State-law timing and strictness used in Part 2	Jurisdiction / firm-year	2003–2023; 51 jurisdictions
Panel B. Mechanism and validation layers			
Filing-level AI-use classification	Customer-facing, internal, third-party, development, and incidental AI labels	Firm-year	2006–2023; 9,322 firm-years / 2,413 firms
Paragraph AI excerpts and direct-text extraction	High-precision adoption indicators and explicit AI-risk links	Firm-year	1994–2023; 86,410 firm-years / 10,661 firms; direct-text exercise on 375 AI-adopting firm-years
Patent-sidecar merge	External check on invention-related activity	Firm-year	1994–2023; 86,410 firm-years / 10,661 firms

Notes. Panel A contains the layers that enter the main disclosure and DBN designs. Panel B contains the narrower mechanism and validation layers. This table is a reader's map, not a full variable dictionary. Pre-2005 breach-risk observations come from voluntary risk-factor disclosures before the SEC mandate; the main disclosure analysis begins in 2005. Detailed variable construction, prompt text, and legal codings are reported in the appendix.

Requiring both non-empty Business section text and non-missing breach-risk attention yields the usable disclosure sample of 51,746 firm-years and 6,980 firms. Adding filing-length controls and lagged assets retains 46,960 firm-years and 6,353 firms, or about 91

per cent of that baseline sample. The fullest lagged-accounting specification retains 18,101 firm-years and 2,908 firms, or 35 per cent. The filing-length and lagged-assets specifications are the main designs, and the full-controls specification is a conservative lower bound. Appendix Table A2 reports the underlying attrition diagnostics.

4.2 Measuring AI Activity

To operationalise the invention-adoption distinction, the paper builds two text-based measures from the Business section. The AI invention measure captures language about models, algorithms, machine learning capability, and related research activity. The AI adoption measure captures language about using AI in products, services, customer interactions, and operating processes. Both are scaled by Business section word count so that they measure AI intensity, not document length. Appendix A.1 reports the full term sets, preprocessing workflow, and construction details.

Although the two series are positively correlated, they are far from identical, which fits a setting in which firms both invent and deploy AI while the two margins remain empirically distinct. Independent validation layers support that distinction. LLM-based AI intensity scores correlate positively with both measures, and corrected AI patent counts provide an external benchmark for invention-related activity. The main adoption measure also removes generic IT language and residual scientific false positives. Appendix construct-validity checks add a separate non-AI digitisation placebo built from e-commerce, cloud services, digital platforms, online channels, mobile apps, data warehouses, enterprise software, SaaS, digital-transformation language, and APIs.

Because the filing-based measures preserve the invention-adoption distinction most cleanly in the main disclosure regressions and in the DBN analysis, they remain the lead variables throughout the paper. For the mechanism analysis, however, the paper also uses a narrower paragraph-based measure. Starting from a high-precision anchor set of unmistakable AI terms, it extracts AI-relevant excerpts from the Business section and classifies them into adoption, invention, customer-facing, internal, third-party, or incidental use. The resulting series is intentionally narrower and substantially sparser than the filing-based measures, so it sharpens the mechanism analysis without replacing the main measures. A human validation exercise (Appendix A.3) confirms 79 per cent case-level agreement on the invention-adoption distinction and 88 per cent on the customer-facing

flag. The broader validation argument is cumulative: the dictionary terms are reported in full, the explicit non-AI digitisation placebo shows that the adoption wedge survives conditioning on broader digital language, and the filing dictionary, paragraph classifier, and direct-text extraction pass all converge on the same deployment-versus-research ordering.

4.3 Measuring Breach-Risk Attention and Supplementary Layers

On the outcome side, breach-risk attention comes from the Risk Factors section. Conceptually, it measures how much of a firm's risk discussion is devoted to unauthorised access, security failure, loss of customer or business data, and related operational consequences. The estimand is disclosed breach-risk salience, not realised breach incidence. For the question studied here, it is useful because it reflects what managers and counsel treat as material enough to disclose.

A 50-topic Latent Dirichlet Allocation model applied to the Risk Factors corpus produces the main outcome, *DB_score*, which is the estimated share of each firm's Risk Factors text devoted to the breach-related topic. The measure captures the prominence of breach risk in the overall disclosure, not just isolated keyword matches. Validation exercises show close alignment with independent LLM-based and dictionary-based alternatives, and appendix checks show qualitative stability across alternative topic counts.

Those filing-level measures still leave two supplementary tasks. The first is to identify use type within adoption. For that purpose, the mechanism analysis adds a filing-level AI-use classification that assigns non-exclusive categories such as customer-facing deployment, internal-use deployment, third-party adoption, development, and incidental mention. The full classification covers 9,322 firm-years over 2006–2023. An overlapping benchmark sample is classified independently, with category-level agreement between 88 and 95 per cent for the substantive labels.

The second task is to recover what firms say when they link AI directly to cyber risk. On the 375 AI-adopting firm-years used for this exercise, 45 contain at least one explicit statement connecting AI or advanced technology to cyber or breach risk. Across 103 directional statements, 101 describe AI as increasing cyber or breach vulnerability and only 2 describe it as reducing that vulnerability. This layer serves as direct evidence on the mechanism, not as a stand-alone source of identification.

Corrected AI patent counts provide an external validation layer, and the filing panel is also merged to state DBN law timing using firms' headquarters states. Supplementary realised-breach data remain in the appendix as a boundary check on scope, not a main-text outcome.

The raw distributions already hint at the same wedge. The average firm-year devotes 3.30 per cent of its risk disclosure to the breach-related topic, but the distribution is wide, with a median of 1.99 per cent and a 75th percentile of 5.19 per cent. Both AI measures are highly skewed, with many firms never mentioning AI and a fast-growing upper tail in later years. The paragraph-based measures confirm that specific AI adoption language remains much sparser than the broader filing-based measures.

Figure 1 plots annual means of the invention measure, the adoption measure, and breach-risk attention, each normalised to 2006 = 100. Two descriptive facts stand out. First, AI adoption and breach-risk attention rise together after 2010, when AI becomes visibly embedded in public-company filings. Second, adoption and invention do not move in lockstep. Invention accelerates more sharply in the final years, whereas adoption begins rising earlier and tracks breach-risk attention more closely. Those descriptive patterns are suggestive, but the next section asks whether the wedge survives joint specifications and richer controls.

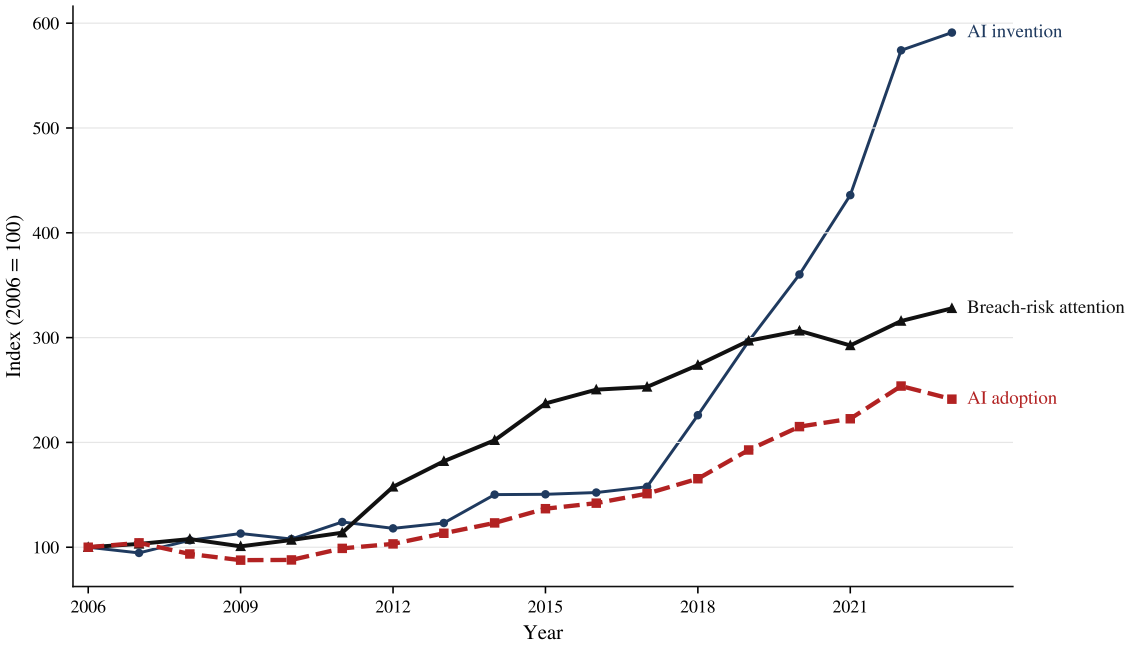

FIGURE 1. AI diffusion and breach-risk attention over time

Notes. The figure plots annual means of the AI invention measure, the AI adoption measure, and breach-risk attention, each normalised to 2006 = 100. The year 2006 is the first year in which both AI series have positive annual means. Annual means are computed from the filing panel with non-empty Business section text and non-missing breach-risk scores.

5. Empirical Strategy

Descriptive patterns alone are not enough, because they do not separate the wedge from common filing trends or compositional differences across firms. To do that, the empirical strategy has two linked parts. The first examines whether firms that describe more AI activity in their business operations also devote more of their risk disclosures to breach concerns. The second examines whether regulation that raises the expected cost of a breach feeds back onto later AI activity. Throughout, the designs are organised around the distinction between AI invention and AI adoption.

5.1 Baseline Disclosure Design

The disclosure analysis estimates a firm-and-year fixed-effects disclosure regression:

$$DB\_score_{it} = \alpha_i + \delta_t + \beta\, AI_{it} + \gamma\, Z_{it} + \varepsilon_{it}, \tag{1}$$

where DB_score_{it} is firm i's breach-risk topic share in year t, AI_{it} is one of the AI measures, α_i are firm fixed effects, and δ_t are year fixed effects. Firm fixed effects absorb time-invariant differences in business model, disclosure style, and sector positioning. Year fixed effects absorb aggregate shifts in AI language, cyber salience, and disclosure conditions.

Only two controls enter the benchmark specification, log Business section word count and log Risk Factors word count. These filing-length controls are essential because both the main regressor and the outcome are constructed from the same filing. They absorb broad disclosure-verbosity shocks without forcing the analysis onto a highly selected sample. Lagged log assets enter next as the main conservative controlled specification, which still retains about 91 per cent of the usable disclosure sample. A fuller lagged accounting vector, including lagged R&D intensity, Tobin's q, tangibility, and return on assets, appears only as a lower-bound sensitivity check. This ordering is deliberate. Many of these firm characteristics are plausibly downstream of AI deployment, and the fullest specification also imposes substantial sample attrition. The benchmark design is the firm-and-year fixed-effects model with filing-length controls, with the assets-only model as the main conservative extension.

For the main disclosure tables, the AI regressors are standardised so coefficients can be read as one-standard-deviation effects. In decomposition exercises where the within-firm and between-firm comparison matters more than direct comparability, the raw AI ratios are used instead.

5.2 Adoption versus Invention

The core specification is a horse-race regression that puts adoption and invention in the same equation:

$$DB\_score_{it} = \alpha_i + \delta_t + \beta_1\, AI_{it}^{\text{adopt}} + \beta_2\, AI_{it}^{\text{invent}} + \gamma\, Z_{it} + \varepsilon_{it}. \tag{2}$$

This specification asks whether breach-risk attention loads on AI activity in general or specifically on the deployment margin. If the relevant channel runs through ongoing data handling, adoption should dominate and invention should be close to zero when both margins enter together. A single composite AI index would blur precisely the distinction the paper is trying to identify, which is why the horse race is central to the empirical design, not a secondary exercise.

Customer-facing deployment is partly a persistent business-model characteristic, so a purely within-firm design could understate economically meaningful variation. The horse race is therefore complemented by Mundlak decompositions that separate each AI measure into firm-level means and within-firm deviations (Mundlak 1978, Wooldridge 2019). If adoption matters in both dimensions while invention remains weak, the case for a generic "technology orientation" explanation becomes harder to sustain. Lead-lag specifications provide additional support on temporal ordering by testing whether changes in AI language predict later changes in breach-risk attention more clearly than the reverse. These are not stand-alone identification designs, but they help discipline the interpretation of direction. Standard errors are clustered at the firm level throughout the disclosure analysis.

5.3 Mechanism and Customer-Facing Deployment

To locate the premium within adoption, the main design compares customer-facing AI firm-years with other substantive AI firm-years using industry-year fixed effects. Customer-facing deployment is mainly a between-firm characteristic of how the business uses AI, not a margin that changes sharply within firms from year to year. Industry-year fixed effects absorb common shocks to technology conditions, product-market environment, and sectoral cyber salience while preserving the cross-sectional variation in deployment type that identifies the channel. A firm-fixed-effects specification at this stage would test a different and much narrower question.

A supporting within-firm design interacts continuous AI-adoption intensity with a customer-facing indicator. It is useful, but it is not the lead mechanism test because firm fixed effects absorb much of the relevant variation in deployment type. The filing-level LLM classification and the paragraph-based measure are used here to sharpen the distinction between customer-facing, internal, third-party, and development-oriented AI use. They

serve as interpretation devices that allow the mechanism to be tested at a finer level of granularity than the dictionary measures alone would permit.

5.4 Supplementary Regulatory Feedback through the DBN Event Study

The supplementary regulatory design tests whether data-breach regulation bears more heavily on deployment than research. It uses the staggered enactment of state Data Breach Notification laws and assigns treatment by the state in which the firm is headquartered, following Miller and Tucker (2009). DBN obligations can also depend on the residence of affected consumers, so headquarters-state assignment is not a perfect measure of treatment exposure. To the extent that customer geography matters beyond headquarters location, the resulting treatment misclassification should attenuate the estimated effect towards zero instead of generating spurious negative results. The main event-study specification is

$$Y_{it} = \alpha_i + \delta_t + \sum_{k \neq -1} \beta_k D_{it}^k + \Gamma X_{i,t-1} + \varepsilon_{it}, \tag{3}$$

where Y_{it} is the AI invention measure or the AI adoption measure, D_{it}^k equals one when firm i is at event time k relative to its home state's DBN enactment, event time -1 is omitted, and $X_{i,t-1}$ contains lagged financial controls.

Because staggered-adoption designs are sensitive to treatment-effect heterogeneity, the paper does not rely on conventional two-way fixed effects alone. It reports the standard TWFE event study together with the imputation estimator of Borusyak, Jaravel, and Spiess (2024) and the group-time approach of Callaway and Sant'Anna (2021). Because treatment varies at the state level, the baseline TWFE design uses state-clustered standard errors.

Identification here is necessarily narrower. The key assumption is that state DBN enactment is unrelated to unobserved changes in firm AI activity that would have occurred anyway. This assumption is plausible because the laws were generally responses to breaches and consumer-protection concerns, not responses to state-level AI activity. Pre-trend diagnostics are cleaner for adoption than for invention, which is one reason the DBN evidence is best read as supportive, not definitive.

Another constraint comes from timing. DBN adoption is front-loaded relative to the period in which AI language becomes common in 10-K filings. The share of untreated firm-years in the estimation sample falls sharply after 2005, so the effective comparison

group is thin in the main AI era. More broadly, the paper studies disclosed breach-risk attention rather than realised breach incidence, and the AI measures capture what firms say about AI in their business descriptions rather than engineering measures of model quality or cyber exposure. The results are therefore best read as evidence on relative emphasis, especially between adoption and invention and within adoption itself.

The disclosure analysis is designed to recover where firms place breach risk inside the business model they describe to investors. The DBN analysis then asks whether a regulatory shock to the cost of data failure changes deployment more than research on the activity side. Read together, the two designs are not trying to identify one universal causal parameter. They ask whether the relative emphasis on deployment over invention appears across multiple empirical margins.

6. Results

6.1 AI Adoption and Disclosed Breach-Risk Attention

The disclosure regressions reveal the wedge immediately. In the benchmark firm-and-year fixed-effects specification with filing-length controls, a one-standard-deviation increase in AI adoption is associated with a statistically significant rise in disclosed breach-risk attention equal to roughly 5 per cent of the sample mean. Panel A of Table 2 and Panel A of Figure 2 show that the coefficient remains positive throughout the specification ladder. Adding filing-length controls narrows the estimate, but it does not overturn it.

Economically, the benchmark coefficient is moderate but meaningful. One standard deviation of the main adoption measure corresponds to about 0.275 additional adoption-term matches per 1,000 Business-section words. In the benchmark joint specification, that shift implies roughly 0.0016 more breach-topic share, or 4.9 per cent of the sample mean. The unconditional standard deviation of breach-risk attention is 0.0376, so the implied shift is about 0.043 standard deviations of the outcome. Relative to the unconditional distribution, it closes about 5 per cent of the distance from the sample median breach-topic share (1.99 per cent) to the 75th percentile (5.19 per cent). The customer-facing premium discussed below is larger, at 0.0088, or 26.8 per cent of the mean and about 27.6 per cent of the median-to-75th-percentile gap.

TABLE 2. Benchmark association and adoption-versus-invention wedge

Measure	Specification	Estimate	Std. error	N	Firms
Panel A. Benchmark specification ladder					
Baseline adoption measure	No controls	0.0025***	0.0006	51,746	6,980
Baseline adoption measure	Filing lengths	0.0021***	0.0005	51,746	6,980
Baseline adoption measure	Filing lengths + assets	0.0015***	0.0005	46,960	6,353
Baseline adoption measure	Full controls	0.0013**	0.0006	18,101	2,908
Preferred adoption measure	No controls	0.0021***	0.0004	51,746	6,980
Preferred adoption measure	Filing lengths	0.0017***	0.0004	51,746	6,980
Preferred adoption measure	Filing lengths + assets	0.0011***	0.0004	46,960	6,353
Preferred adoption measure	Full controls	0.0010**	0.0005	18,101	2,908
Panel B. Adoption-versus-invention horse races					
Baseline text measure	Adoption	0.0025***	0.0006	51,698	6,973
Baseline text measure	Invention	0.0001	0.0004	51,698	6,973
Preferred text measure	Adoption	0.0021***	0.0004	51,698	6,973
Preferred text measure	Invention	0.0000	0.0005	51,698	6,973
Paragraph measure	Adoption, full sample	0.0008**	0.0003	51,723	6,974
Paragraph measure	Invention, full sample	0.0003	0.0002	51,723	6,974
Paragraph measure	Adoption, 2012–2023	0.0009**	0.0004	34,251	4,939
Paragraph measure	Invention, 2012–2023	0.0005**	0.0002	34,251	4,939

Notes. Panel A shows the benchmark ladder for the baseline and preferred adoption measures. Panel B shows horse-race comparisons between adoption and invention using the baseline text measure, the preferred text measure, and the paragraph measure. The preferred text measure combines the unchanged pruned invention dictionary with the preferred adoption dictionary. Significance levels are * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Reading down the control ladder is most useful as an economic exercise, not a mechanical one. Adding lagged assets leaves the estimate clearly positive while retaining about 91 per cent of the usable disclosure sample. The full-controls specification cuts the sample to roughly one third and is better read as a lower bound from a compositionally selected sample than as the benchmark estimate. The main adoption measure follows the same downward slope through the ladder but remains positive and statistically significant throughout.

It becomes clearest when adoption and invention enter together. Panel B of Table 2 reports the horse race, and Panel B of Figure 2 makes the wedge visible. Once both mar-

gins enter simultaneously, adoption remains positive and statistically strong in both the baseline and benchmark filing-level text measures, whereas invention is economically negligible and statistically indistinguishable from zero in those lead filing-level measures. The key contrast is between firms that describe AI as an operating technology embedded in products and services and firms that describe AI as a developmental or capability-building activity. On that margin, adoption is more tightly associated with recurrent data handling and with the breach-risk salience that accompanies it.

TABLE 3. Within-between decomposition of raw AI ratios

Component	Simple adoption	Full-controls adoption	Horse race
Panel A. Adoption components			
Within-firm adoption ratio	7.713*** (1.663)	4.287** (1.669)	7.025*** (1.762)
Between-firm adoption ratio	12.482*** (2.339)	12.832*** (3.210)	11.810*** (2.430)
Panel B. Invention components			
Within-firm invention ratio	–	–	1.560 (1.236)
Between-firm invention ratio	–	–	2.412 (1.608)
Panel C. Observations			
N	51,745	18,100	51,745

Notes. Entries report coefficient and standard error in parentheses. The Mundlak models use the raw AI ratios rather than the standardised headline regressors, so the magnitudes are not directly comparable to Table 2. The simple and horse-race columns use the retained 1997–2023 decomposition panel, and the full-controls column uses the retained 1998–2023 panel, so each column is one observation smaller than the corresponding benchmark specification. The full-controls column also adds the lagged accounting vector and firm means. Significance levels are * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

The most direct construct-validity concern is that the adoption dictionary could be proxying for broad digitisation rather than AI-specific deployment. Appendix C.1 therefore adds an explicit non-AI digitisation placebo built from e-commerce, cloud services, digital platforms, mobile apps, SaaS, APIs, data warehouses, enterprise software, and digital-transformation language. When that placebo enters alongside adoption and invention, the adoption coefficient remains positive at 0.0016 (standard error 0.0004), the invention

coefficient remains near zero in the lead filing-level specification, and the placebo itself is also positive. The placebo coefficient is somewhat larger because it captures a broader bundle of digital-transformation language with more textual mass than the narrower AI-deployment dictionary. Broad digitisation language is therefore relevant, but it does not subsume the adoption wedge. Appendix C.1 reaches the same conclusion from the topic side: once filing-length controls are added, the number of positive BH-significant topics falls from 11 to 2 while Topic 45 remains the top positive loading throughout.

The joint specification also addresses the same-filing confound directly. If the result were driven entirely by a filing-production process that simultaneously expanded technology language in the Business section and cyber language in the Risk Factors section, one would expect both the adoption and invention coefficients to rise. They do not. Invention remains flat while adoption persists. Appendix diagnostics that remove shared vocabulary, add the non-AI digitisation placebo, and substitute external patent measures leave the sign pattern unchanged. The narrower paragraph-based measure is useful for the mechanism analysis but compresses the wedge in later years, partly because later invention language itself appears more data-intensive, so it remains a supporting layer rather than the lead filing-level disclosure measure.

The wedge is not purely cross-sectional, either. Table 3 shows that adoption loads positively both within firms and between firms, with the between-firm component larger. This pattern fits deployment type as partly a persistent business-model characteristic, but it is also consistent with already-exposed firms selecting into adoption. At the same time, the within-firm adoption effect is clearly positive, so the result is not merely a static sort across firms. In the horse-race Mundlak model, adoption remains significant in both dimensions while invention remains insignificant in both. The decomposition lines up with Prediction 1 while still leaving room for non-random selection into the adoption margin.

Even so, the estimand remains *disclosed breach-risk attention*, not realised breach incidence. The coefficient does not imply that adoption mechanically raises realised breach incidence in every instance, nor that invention is irrelevant for governance. It shows instead that disclosed breach-risk attention is more concentrated where firms describe AI as part of an operating data process. Appendix timing checks are directionally aligned with adoption preceding later breach-risk attention, which is consistent with the main reading.

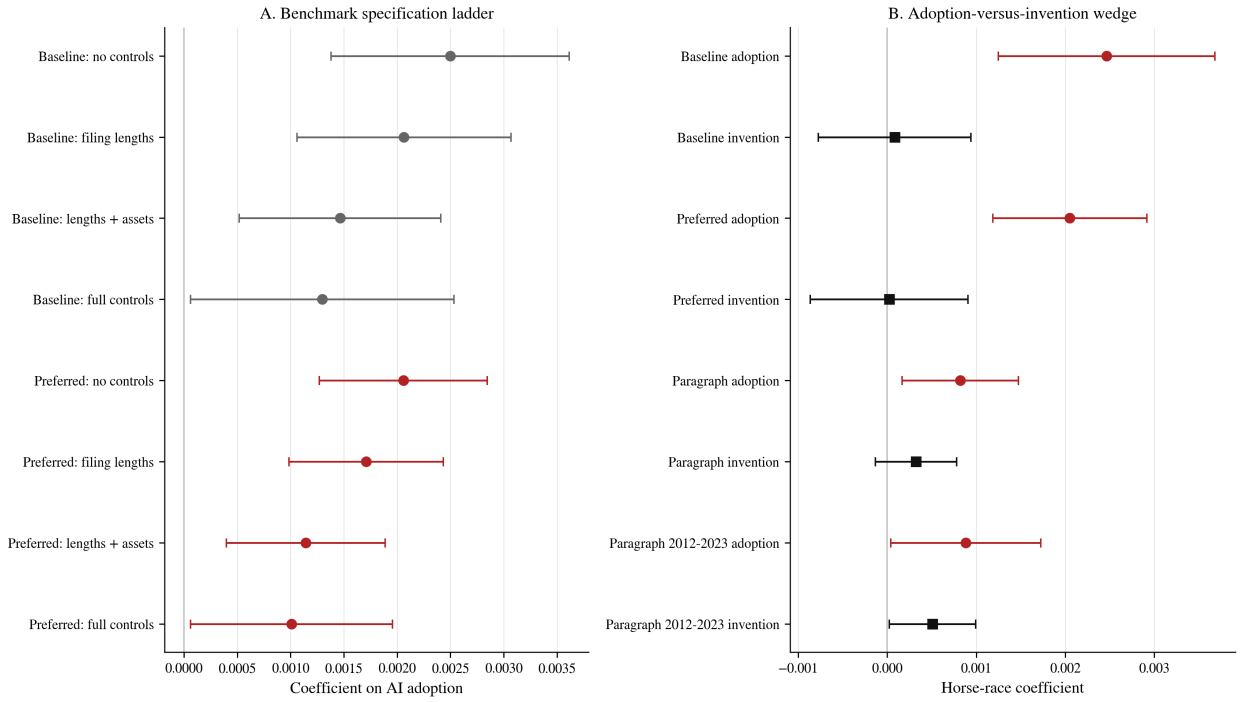

FIGURE 2. Specification ladder and adoption-versus-invention wedge

Notes. Panel A plots the AI adoption coefficient across the baseline and preferred text measures under four control specifications with no controls, filing lengths, filing lengths plus lagged assets, and full controls. All specifications include firm and year fixed effects. Panel B plots the horse-race coefficients for adoption and invention under three measurement approaches, namely the baseline text measure, the preferred text measure, and the paragraph measure (full sample and post-2012 subsample). Whiskers show 95 per cent confidence intervals.

6.2 The Customer-Facing Mechanism

Table 4 shows that the premium sits in customer-facing deployment. Prediction 2 concerns a business-model characteristic that is mainly between firms, not within firms, so the lead design compares customer-facing AI firm-years with other substantive AI firm-years within industry-year cells.

In the filing-level LLM classification sample, customer-facing AI firm-years carry a breach-risk premium of 0.0113 relative to other substantive AI firm-years within industry-year cells. The paragraph-based binary mechanism, estimated on the stricter subset of firm-years with substantive adoption excerpts, produces an even sharper premium of 0.0139. When the explicit non-AI digitisation placebo is added to the paragraph design, the

customer-facing premium remains positive at 0.0088 ($p = 0.022$). Not all AI deployment is equivalent. The part most tightly associated with disclosed breach-risk attention is the part embedded in direct interaction with customers, their transactions, and their data.

This evidence is mainly cross-sectional by construction. Customer-facing deployment is largely a persistent business-model characteristic, not a margin that reassigns sharply within firms year by year. The customer-facing premium survives a demanding NAICS4-year fixed-effect specification at 0.0094 ($p = 0.017$). NAICS4-year cells are the main partition because they track current product-market boundaries more closely while preserving more usable within-cell variation for this design. The evidence therefore speaks most clearly to a cross-sectional business-model pattern rather than to a within-firm causal isolate.

The underlying text supports the same interpretation. Of 375 AI-adopting firm-years, 45 contain at least one explicit statement linking AI or advanced technology to cyber or breach risk. Across the 103 directional statements, 101 describe AI as increasing risk and only 2 describe it as reducing risk. An exact binomial test overwhelmingly rejects a symmetric split. This directional imbalance is difficult to reconcile with a pure boilerplate interpretation. If the regressions merely captured generic co-movement between technology language and security language, one would not expect such a pronounced skew in the underlying text. The content of the statements matches the data-handling channel directly, even if it does not by itself adjudicate between customer-facing and non-customer-facing firms.

The supporting checks do not overturn that reading. Appendix C.1 shows that once filing-length controls are added, the same-filing pattern narrows sharply and the breach topic remains the dominant positive loading in the 50-topic specificity test. Table A11 and the other appendix checks show that tail treatments, extensive-margin replacements, alternative breach measures, and patent-based AI measures all preserve the sign pattern. Those exercises are useful checks, but the main-text argument still rests on the wedge, the mechanism, and the supplementary regulatory check that follows.

TABLE 4. Customer-facing mechanism and direct textual evidence

Evidence layer	Value	SE / share	N
Panel A. Customer-facing premium			
Filing-level classification	0.0113***	0.0028	1,966
Paragraph classification	0.0139***	0.0034	994
Paragraph + NAICS4-year FE	0.0094**	0.0039	1,016
Paragraph + SIC-year FE	0.0050	0.0049	1,016
Paragraph + NAICS4-year FE + digitisation placebo	0.0088**	0.0038	1,016
Panel B. Within-firm slope comparison			
Filing-level: customer-facing firms	0.0009*	0.0005	4,986
Filing-level: other AI firms	0.0008	0.0019	4,986
Filing-level equality-test p-value	0.97	p-value	4,986
Paragraph: customer-facing firms	0.0002	0.0006	17,578
Paragraph: other AI firms	-0.0005	0.0005	17,578
Paragraph equality-test p-value	0.37	p-value	17,578
Panel C. Direct text			
Firm-years with explicit AI-risk link	45 / 375	12.0%	375
Statements saying AI raises risk	101		
Statements saying AI reduces risk	2		

Notes. Panel A reports the customer-facing premium under the filing-level binary design, the paragraph binary design, finer industry-year partitions, and the paragraph design with the explicit non-AI digitisation placebo. Panel B reports supporting within-firm slope comparisons and equality-test p-values. Panel C reports evidence from firms' own risk statements. Significance levels are * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

6.3 Supplementary Regulatory Evidence from DBN Laws and AI Diffusion

The DBN design is best read as a supplementary regulatory check rather than as a co-equal design. The prediction is that a legal shock raising the expected cost of a failed data process should align more closely with deployment than with research. The estimated signs follow that ranking. Appendix B and Appendix C.3 report the full tables, figures, and jackknife diagnostics. The main-text reading is therefore narrow: the regulatory evidence reinforces

the disclosure results but is not intended to stand alone as an independent causal design.

7. Conclusion

The results point to a clear conclusion. The liability side of AI diffusion shows up most clearly in deployment, especially customer-facing deployment, rather than in invention. Firms that describe AI as part of a live operating process devote more disclosure space to breach-related concerns, and firms' own risk statements describe AI as expanding exposure far more often than reducing it. The adoption association survives an explicit control for non-AI digitisation language in the same filing text. The DBN evidence is more limited, but its signs are consistent with the disclosure results.

This result matters because public debate still tends to bundle AI research, AI deployment, privacy, and consumer protection into one broad question about whether regulation helps or hinders innovation. The evidence here suggests a narrower and more useful framing. The main association is concentrated in systems that handle personal data on a continuing basis. Consent rules, data-minimisation requirements, vendor oversight, security controls, and breach-notification obligations are therefore more tightly matched to customer-facing deployment than to broad restrictions on AI research and development. The invention-versus-adoption contrast is useful on this point because it shows where disclosed breach-risk salience concentrates inside firms' stated business models.

The managerial reading is just as direct. Deploying AI is not only a decision about expected productivity gains. It is also a decision about whether the firm can govern the data flows that make those gains possible. Customer-facing systems bundle value creation with security, compliance, vendor management, and incident-response obligations. Firms that underestimate those complements may overstate the net value of deployment even when the underlying technology is strong. For investors, too, AI language in filings should not be read only as a signal of growth opportunity. The type of AI activity matters for risk as well as return.

The claims nonetheless remain bounded. The main outcome is disclosed breach-risk attention, not realised breach incidence, so the evidence bears most directly on managerial salience and revealed materiality. Firm and year fixed effects do not rule out selection of already-exposed firms into AI adoption, and the large between-firm component is consis-

tent with that possibility. The filing-based adoption measure is supported by convergent diagnostics across the dictionary, paragraph, and direct-text layers, and the DBN evidence remains supplementary rather than definitive. Those limits narrow the claims. They do not overturn the central result. In the AI economy, data is an asset only under governance conditions that keep it from becoming a liability.

References

Acemoglu, D., D. Autor, J. Hazell, and P Restrepo. 2022. "Artificial intelligence and jobs: Evidence from online vacancies." Journal of Labor Economics, 40(S1):S293-S340.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2019. "Economic Policy for Artificial Intelligence." *Innovation Policy and the Economy* 19: 139–159.

Akey, Pat, Stefan Lewellen, Inessa Liskovich, and Christopher Schiller. 2024. "Hacking Corporate Reputations." Working Paper.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics* 151: 103745.

Bao, Yang, and Anindya Datta. 2014. "Simultaneously Discovering and Quantifying Risk Types from Textual Risk Disclosures." *Management Science* 60 (6): 1371–1391.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3: 993–1022.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *Review of Economic Studies* 91 (6): 3253–3285.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2023. "Generative AI at Work." Working Paper 31161, National Bureau of Economic Research.

Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu. 2024. "Business News and Business Cycles." *Journal of Finance* 79 (5): 3105–3147.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.

Campbell, John L., Hsinchun Chen, Dan S. Dhaliwal, Hsin min Lu, and Logan B. Steele. 2014. "The Information Content of Mandatory Risk Factor Disclosures in Corporate Filings." *Review of Accounting Studies* 19 (1): 396–455.

Dye, Ronald A. 1985. "Disclosure of Nonproprietary Information." *Journal of Accounting Research* 23 (1): 123–145.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (12): 2195–2217.

Florackis, Christos, Christodoulos Louca, Roni Michaely, and Michael Weber. 2023. "Cybersecurity Risk." *Review of Financial Studies* 36 (1): 351–407.

Foerderer, Jens, and Sebastian Schuetz. 2022. "Data Breach Announcements and Stock Market Reactions: A Matter of Timing?" *Management Science* 68 (10): 7298–7322.

Frey, Carl Benedikt, Giorgio Presidente, and Pablo Andres. 2024. "Redirecting AI: Privacy Regulation and the Future of Artificial Intelligence." CEPR Discussion Paper.

Gerlicher, Amelia M. 2023. "Security Breach Notification Chart." Perkins Coie.

Goldfarb, Avi, and Catherine Tucker. 2012. "Privacy and Innovation." In *Innovation Policy and the Economy*, vol. 12, 65–90.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2019. "Firm-Level Political Risk: Measurement and Effects." *Quarterly Journal of Economics* 134 (4): 2135–2202.

Jamilov, Rustam, Helene Rey, and Ahmed Tahoun. 2025. "The Anatomy of Cyber Risk." Journal of Financial Economics, forthcoming.

Kamiya, Shinichi, Jun-Koo Kang, Jungmin Kim, Andreas Milidonis, and Rene M. Stulz. 2021. "What Is the Impact of Successful Cyberattacks on Target Firms?" *Journal of Financial Economics* 141 (2): 411–437.

McLemore, Peyton, and Vasil Mihov. 2026. "Artificial Intelligence and Operational Risk." Review of Corporate Finance Studies, forthcoming.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. "Distributed Representations of Words and Phrases and Their Compositionality." *Advances in Neural Information Processing Systems* 26.

Miller, Amalia R., and Catherine Tucker. 2009. "Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records." *Management Science* 55 (7): 1077–1093.

Mundlak, Yair. 1978. "On the Pooling of Time Series and Cross Section Data." *Econometrica* 46 (1): 69–85.

Rishabh, Krishna, Roxana Mihet, and Julian Jang-Jaccard. 2025. "Cyber Risk and Artificial Intelligence Innovation." Working Paper.

Sautner, Zacharias, Laurence van Lent, Grigory Vilkov, and Ruishen Zhang. 2023. "Firm-Level Climate Change Exposure." *Journal of Finance* 78 (3): 1449–1498.

Verrecchia, Robert E. 1983. "Discretionary Disclosure." *Journal of Accounting and Economics* 5: 179–194.

Wooldridge, Jeffrey M. 2019. "Correlated Random Effects Models with Unbalanced Panels." *Journal of Econometrics* 211 (1): 137–150.

A. Additional Data, Measurement, and Sample Construction

This appendix documents four parts of the study's data and measurement framework: the variable guide and source summary, the sample-construction ladder and attrition diagnostics, the detailed construction of the filing-based and paragraph-based measures, and the treatment-support features of the DBN law data.

TABLE A1. Variable guide and source summary

Variable name	Label in manuscript	Source layer	Construction	Use in analysis	Coverage
AI1 Ratio	AI invention measure	Business section of the 10-K (Item 1)	Matched invention terms divided by Business section word count	Main invention measure in disclosure and DBN analysis	84,012 firm-years / 10,456 firms
AI2 Ratio	AI adoption measure	Business section of the 10-K (Item 1)	Matched adoption terms divided by Business section word count	Main adoption measure in disclosure and DBN analysis	84,012 firm-years / 10,456 firms
DB Score	Breach-risk attention	Risk Factors section of the 10-K (Item 1A)	LDA topic-45 share of Risk Factors text	Main disclosure outcome	51,771 firm-years / 6,981 firms
Paragraph Adoption Any	Paragraph adoption indicator	Paragraph-level LLM classification of Business section anchor excerpts	Binary indicator equal to one if any adoption-classified paragraph appears in the firm-year	Mechanism and validation layer	86,410 firm-years / 10,661 firms
Paragraph Customer-Facing Any	Paragraph customer-facing indicator	Paragraph-level LLM classification of Business section anchor excerpts	Binary indicator equal to one if any customer-facing AI paragraph appears in the firm-year	Customer-facing mechanism layer	86,410 firm-years / 10,661 firms
AI Intensity	Filing-level AI intensity	Filing-level LLM classification output	0–10 scalar AI intensity score from the filing-level classification task	Validation and descriptive support for filing-level classification	9,322 firm-years / 2,413 firms
Corrected AI Patent Count	Corrected AI patent count	Patent-sidecar merge	Corrected raw patent count at the firm-year level	External invention validation	86,410 firm-years / 10,661 firms
Had Breach	Realised breach indicator	Matched PRC and VCDB breach-event panels	Binary indicator equal to one if any realised breach is linked to the firm-year	Appendix boundary check	86,410 firm-years / 10,661 firms
Industry AIOE	Industry AI occupational exposure	Industry AI occupational exposure merge	Industry-level AI occupational exposure score	Industry-level control and IV ingredient	86,407 firm-years / 10,659 firms
Strictness Index	DBN strictness index	Perkins Coie state law coding	0–1 composite across six law dimensions	DBN dose-response support	83,392 firm-years / 10,289 firms
Lagged Log Assets	Lagged log assets	CRSP-Compustat accounting merge	One-period lag of log total assets	Preferred conservative control	74,254 firm-years / 9,059 firms
Lagged Log R&D	Lagged log R&D	CRSP-Compustat accounting merge	One-period lag of log R&D expenditure	Full-controls lower-bound specification	28,071 firm-years / 4,026 firms
Lagged Tobin's q	Lagged Tobin's q	CRSP-Compustat accounting merge	One-period lag of Tobin's q	Full-controls lower-bound specification	74,254 firm-years / 9,059 firms
Lagged Tangibility	Lagged tangibility	CRSP-Compustat accounting merge	One-period lag of PP&E over assets	Full-controls lower-bound specification	70,250 firm-years / 8,756 firms
Lagged Return on Assets	Lagged return on assets	CRSP-Compustat accounting merge	One-period lag of income over assets	Full-controls lower-bound specification	74,254 firm-years / 9,059 firms

Notes. This appendix table is a reference guide. It records where each variable comes from, how it is constructed, and where it is used in the analysis. Technical variable names are retained here for replication convenience, while the main text uses substantive labels.

TABLE A2. Sample construction and attrition diagnostics

Panel A: Sample ladder

Stage	Firm-years	Firms	Share of full panel	Share of usable sample
Full filing panel	86,410	10,661	100.0%	–
Business-section text panel	84,012	10,456	97.2%	–
Risk Factors text panel	51,771	6,981	59.9%	–
Usable disclosure sample	51,746	6,980	59.9%	100.0%
Preferred controlled sample	46,960	6,353	54.3%	90.8%
Full-controls lower bound	18,101	2,908	20.9%	35.0%

Panel B: Full-controls attrition relative to the usable disclosure sample

Variable	Kept mean	Dropped mean	Std. diff.	Comment
Breach risk attention	0.0318	0.0336	-0.047	Similar means.
AI adoption measure ratio	0.0001	0.0001	0.176	Kept sample is more adoption-intensive.
AI invention measure ratio	0.0001	0.0000	0.194	Kept sample is more invention-intensive.
Lagged log assets	6.1072	7.4721	-0.666	Kept sample is smaller.
Lagged Tobin's q	2.6349	1.5606	0.569	Kept sample has higher q.
Data-intensive industry	0.1678	0.5432	-0.853	Dropped sample is more data-intensive.
Technology-industry indicator	0.2078	0.0448	0.506	Kept sample is more tech-oriented.

Notes. Panel A makes the control ladder explicit. The full-controls specification keeps only 35.0% of the usable disclosure sample. Panel B compares the retained and dropped observations using the existing attrition diagnostic output.

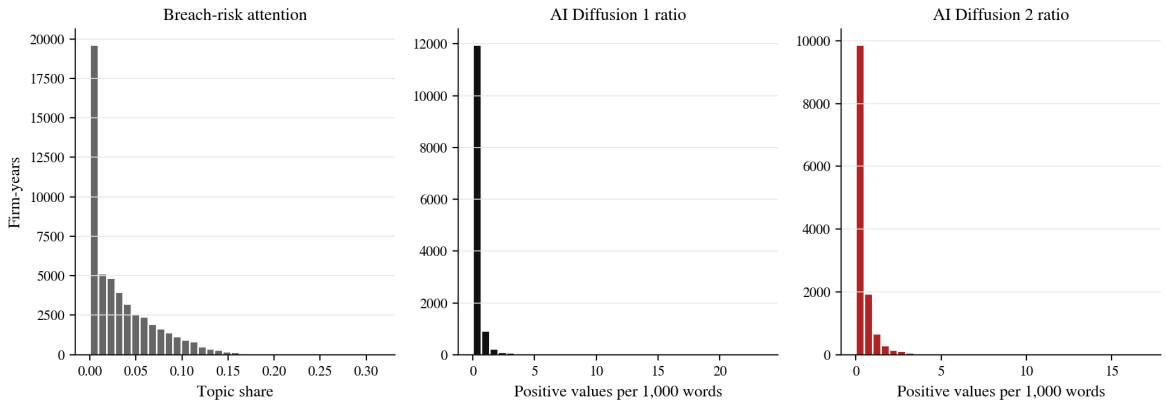

FIGURE A1. Core measure distributions

Notes. The breach-risk distribution is broad. The median topic share is 0.0199 and the 75th percentile is 0.0519. The AI measures are zero-inflated, with zero shares of 84.0% for the AI invention measure and 84.2% for the AI adoption measure. The AI histograms therefore plot positive observations only, scaled to mentions per 1,000 Business-section words, so the upper tail remains visible.

A.1 Text Processing, Dictionary Construction, and the LDA Outcome

The main text describes the construction of the AI invention measure, the AI adoption measure, and breach-risk attention in Section 4.2. This subsection records the preprocessing workflow, the invention and adoption term sets, the Word2Vec expansion objective, the LDA likelihood, the full topic visualisations, and the alternative topic-count checks that are not shown in the main text.

All filing-based text measures begin from the same preprocessing pipeline. Business section text and Risk Factors text are lower-cased, tokenised, stripped of punctuation and numerical symbols, filtered for tokens shorter than three characters, cleaned of common stop words, and Porter-stemmed. As an illustration, the sentence "The single fleet operating models, anchored by Boeing MAX aircraft for the Mainline fleet and Embraer E175 aircraft for the Regional fleet, will drive more productivity and cost efficiency in the business" from Alaska Air Group's 2023 filing becomes, after stemming, `singl fleet oper model anchor boe max aircraft mainlin fleet embraer aircraft region fleet drive product cost effici busi`.

Tables A3 and A4 record the invention term set and the preferred adoption term set used in the paper.

TABLE A3. AI invention term set ('.' denotes any characters)

Term 1	Term 2	Term 3	Term 4
abstract.strategy.games	adaboost	artificial.intelligen	auto.encod
back.propagat	bayesian.network	convolutional.neural.net	computer.vision
data.mining	decision.tree	deep.belief.net	deep.learn
deep.neural.net	deep.reinforcement	dimensional.reduc	ensemble.learn
feature.engineer	feature.extract	feature.select	fuzzy.environment
fuzzy.logic	fuzzy.number	fuzzy.set	fuzzy.system
fuzzy-c	gaussian.model	gaussian.process	genetic.algorithm
genetic.program	gradient.descent	gradient.tree.boost	high.dimensional.data
high.dimensional.feature	image.generation	image.recognition	kernel.learn
k-means	language.modeling	language.process	logistic.regression
machine.intelligen	machine.learn	map.reduce	memet.algorithm
multilayer.perceptron	natural.language.process	natural.gradient	neural.network
neural.turing	object.recogni	particle.swarm	pattern.recogni
random.forest	rankboost	recurrent.neural.net	reinforcement.learn
restricted.boltzmann	sentiment.analy	sparse.code	sparse.represent
spectral.cluster	speech.recognition	stochastic.gradient	supervised.learn
support.vector.machine	transfer.learn	translation	vector.machine
visual.question.answering	AI		

Notes: Entries are the 70 unique regex patterns used to construct the AI invention measure. The symbol '.' denotes any intervening characters, so `deep.belief.net` and `restricted.boltzmann` are single patterns rather than separate tokens.

TABLE A4. Preferred AI adoption term set (underscores denote spaces after stemming)

Term 1	Term 2	Term 3	Term 4
adapt_learn	advanc_analyt	algorithm_analyz	algorithm_develop
analyt_platform	artifici_intellig	automat_detect	big_data
captur_data	charg_captur	clover_assist	comput_intens
control_logic	cut_edg	data_analysi	data_analyt
data_manipul	data_mine	data_model	data_scienc
databas_data	decis_tree	deep_learn	digit_workflow
direct_interfac	dispar_data	extract_data	faster_accur
high_accur	input_data	insight_analyt	insight_data
intellig_data	intellig_machin	intellig_platform	intellig_technolog
internet_thing	machin_learn	natur_languag	neural_network
onlin_analyt	pattern_recognit	predict_analyt	predict_model
process_dsp	realtim_analyt	rich_data	robot_process
sensor_fusion	simpl_easi	singl_view	solid_model
sophist_algorithm	techniqu_identifi	visual_tool	

Notes. Entries are the stemmed patterns used to construct the preferred AI adoption measure after removing residual biomedical and scientific false positives. Underscores denote spaces after stemming. Table A9 reports the additional overlap and generic-IT diagnostics for this measure.

The Word2Vec expansion step uses the Skip-Gram variant (Mikolov et al. 2013). Given a corpus $\{w_1, \ldots, w_T\}$, the model maximises $\sum_{t=1}^{T} \sum_{-j \le c \le j, c \neq 0} \log p(w_{t+c} \mid w_t)$, where j is the context window and the conditional probability is $p(w_o \mid w_c) = \exp(u_o v_c) / \sum_{w=1}^{W} \exp(u_w v_c)$, with W the vocabulary size. The resulting representations are used to identify AI-related neighbouring phrases in the business-description corpus for the construction of the AI adoption measure.

The breach-risk outcome is built from the Risk Factors section using Latent Dirichlet Allocation (Blei, Ng, and Jordan 2003). Each document is treated as a mixture of latent

topics, with the corpus likelihood

$$p(D \mid \alpha, \beta) = \prod_{d=1}^{D} p(\theta_d \mid \alpha) \prod_{n=1}^{N_d} \sum_{z_{d,n}} p(z_{d,n} \mid \theta_d)\, p(w_{d,n} \mid z_{d,n}, \beta).$$

The model is estimated with Gibbs sampling. The preferred model has 50 topics. Topic 45 is the breach-related topic, and breach-risk attention is the share of Risk Factors text allocated to it. Figures A2 and A3 display all 50 topics from the preferred model. Alternative topic counts leave the same substantive conclusion. At $K = 30$, the selected breach-like topic is the top positive loading, only 2 positive topics survive BH correction, and the horse race leaves adoption positive at 0.002177 while invention is essentially zero. At $K = 100$, breach content splits across adjacent cyber topics, but the adoption-versus-invention sign pattern remains intact.

FIGURE A2. Topics 0–24 generated by the LDA model

FIGURE A3. Topics 25–49 generated by the LDA model

A.2 LLM Classification and Paragraph Measures

The main text introduces the filing-level and paragraph-level LLM layers. This subsection records the sample composition, sparsity, and validation role of those layers without repeating the main mechanism argument.

The filing-level use-type classification assigns non-exclusive categories (development,

third-party adoption, customer-facing deployment, internal-use deployment, incidental mention, and no meaningful AI use) to 9,322 firm-years with non-zero AI text. The paragraph-based measure starts from high-precision AI-anchor paragraphs in the Business section, classifies anchor-centred excerpts into adoption, invention, customer-facing, internal, third-party, or incidental use, and aggregates those excerpt-level classifications to the firm-year level. Table A5 summarises both layers. The paragraph measures are sparse by design. Only a small minority of firm-years contain sufficiently specific AI text to support high-precision paragraph coding. Figure A4 makes the trade-off visible.

TABLE A5. LLM classification and paragraph-measure summary

Panel A: Filing-level LLM classification sample		
Metric	Count / value	Share of classified sample
Customer-facing AI	2,237	24.0%
Internal-use AI	1,370	14.7%
AI development	1,562	16.8%
Third-party adoption	1,207	12.9%
Incidental mention	1,255	13.5%
No meaningful AI use	4,800	51.5%
Mean AI intensity	1.6500	
Panel B: Paragraph-based measurement layer		
Metric	Value	Type
Paragraph adoption-any share	2.4%	Share
Paragraph invention-any share	0.5%	Share
Paragraph customer-facing-any share	1.7%	Share
Paragraph adoption-any share in 2023	11.7%	Share
Pearson corr. with filing-level AI adoption measure	0.3055	Correlation
Pearson corr. with filing-level AI invention measure	0.2736	Correlation
Pearson corr. with corrected AI patents	0.0483	Correlation
Pearson corr. with LLM AI intensity	0.2894	Correlation
Spearman corr. of LLM breach score with breach-risk attention	0.7014	Correlation

Notes. The filing-level use-type categories are non-exclusive, so the category shares in Panel A do not sum to 100 per cent. The paragraph measure is intentionally sparse and is reported here to clarify why it is used as a high-precision mechanism layer, not as the lead paper-wide measure.

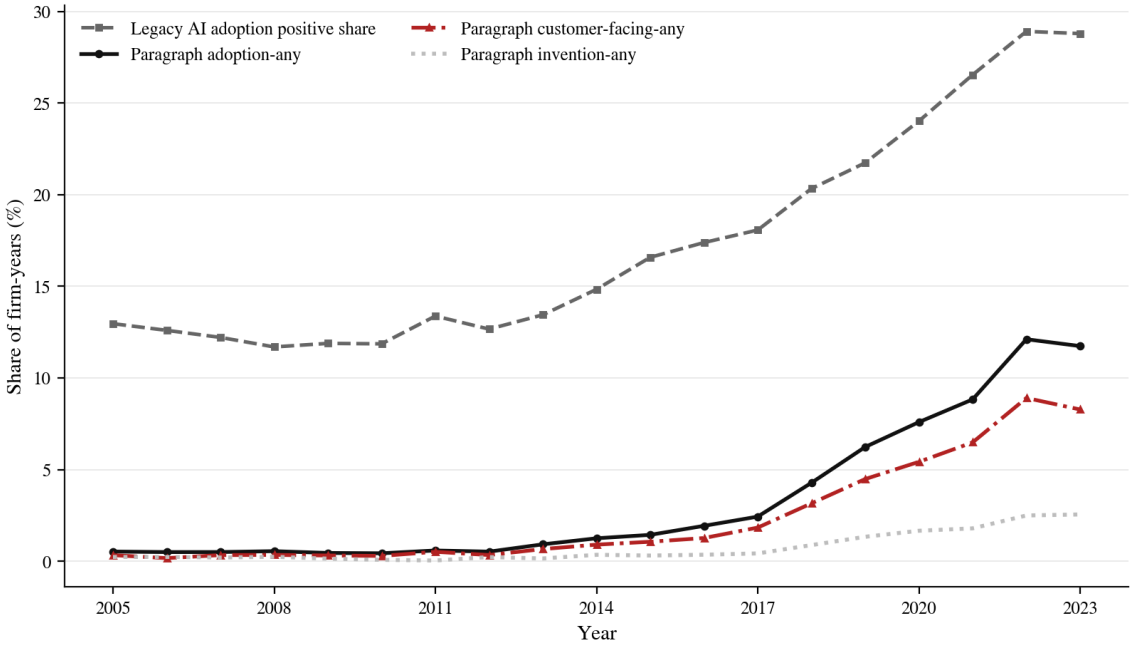

FIGURE A4. Evolution of the paragraph-based AI measures

Notes. The figure plots the share of firm-years with positive values for four AI measures. By 2023, the filing-level AI adoption measure is positive in 28.8% of firm-years, while the paragraph adoption-any share reaches 11.7% and the paragraph customer-facing share reaches 8.3%. The figure makes the measurement trade-off visible because the paragraph series tracks the same broad rise in AI language but keeps a much narrower signal.

A.3 Human Validation of the Invention-Adoption Classification

To test whether the paragraph-level classification reflects the underlying text, 100 firm-years (25 per dictionary-assignment stratum) are hand-coded by an independent coder working from the selected AI-relevant excerpts, with no access to the machine labels. The human coder agrees with the paragraph classification on 79 of 100 cases for the primary orientation label and 88 of 100 for the customer-facing flag. The 21 disagreements are structured. Twelve arise from conservative compression (mixed cases reduced to one side, or thin-signal excerpts downgraded to neither) and 9 from boundary disputes where development and deployment language co-occur in the same excerpt set. The case-level comparison, confusion matrix, and disagreement diagnoses are documented in the replication package.

No separate hand-coded gold standard is reported for the broad filing-level dictionary because that measure is designed as a full-filing count rather than an excerpt-level classifier. If the remaining disagreement is mostly classical measurement error in the filing-level regressor, it should attenuate adoption coefficients towards zero rather than create spurious positives. The fact that the filing dictionary, the narrower paragraph classifier, and

the direct-text extraction layer all point to the same invention-adoption ranking therefore makes the main coefficient more likely to be conservative than inflated.

A.4 DBN Timing, Strictness, and Treatment Structure

The DBN data layer is built from the Perkins Coie Security Breach Notification Chart (Gerlicher 2023), a standardised legal reference that summarises each jurisdiction's breach-notification statute. For the 52 jurisdictions, the coding tracks six dimensions:

a. **Notification deadline.** The number of days within which firms must notify affected individuals after discovering a breach. States without a numerical deadline are assigned the sample maximum before inversion.
b. **Penalty provisions.** Whether the statute specifies financial penalties for non-compliance and, when it does, the maximum penalty amount.
c. **Private right of action.** Whether affected individuals can bring civil suits for non-compliance.
d. **Scope of covered data.** A three-point scale that codes whether the statute uses a narrow, medium, or broad definition of covered personal information.
e. **Encryption safe harbour.** Whether notification is excused when the breached data was encrypted.
f. **Attorney General notification.** Whether separate notification to the state Attorney General is required.

Each dimension is normalised to the $[0, 1]$ interval, with higher values indicating stricter regulation. Notification deadlines are inverted so that shorter deadlines are stricter; maximum penalties are scaled by the sample maximum; private action and Attorney General notification are binary; scope is coded as 0, 0.5, or 1; and encryption safe harbours enter negatively because they reduce effective strictness. The composite index is the simple average across the six normalised dimensions. A ten-state audit against the underlying statute text was conducted for California, New York, Texas, Florida, Illinois, Massachusetts, Ohio, Georgia, Virginia, and Alabama and confirmed the coding on the audited dimensions.

Table A6 reports the enforcement date, strictness index, notification deadline, private right of action, and Attorney General notification requirement for each jurisdiction. Figure A5 shows why the DBN design must be interpreted with care. The untreated share of firm-years collapses quickly after the early adoption wave, falling from 58.4% in 2005 to 0.9% in 2015 and zero by 2023.

TABLE A6. DBN law timing and strictness by jurisdiction

State	State name	Date	Strictness	Deadline	Private action	AG notice
CA	California	2003-07-01	0.6667	30 days	Yes	Yes
GA	Georgia	2005-05-05	0.0833	reasonable delay	No	No

Table A6 continued

State	State name	Date	Strictness	Deadline	Private action	AG notice
ND	North Dakota	2005-06-01	0.2500	reasonable delay	No	Yes
DE	Delaware	2005-06-28	0.5833	60 days	No	Yes
TN	Tennessee	2005-07-01	0.4167	not specified	Yes	No
WA	Washington	2005-07-24	0.5000	reasonable delay	Yes	Yes
AR	Arkansas	2005-08-12	0.3333	reasonable delay	No	Yes
NV	Nevada	2005-10-01	0.2500	reasonable delay	No	Yes
NC	North Carolina	2005-12-01	0.5833	reasonable delay	Yes	Yes
NY	New York	2005-12-07	0.4167	30 days	No	Yes
CT	Connecticut	2006-01-01	0.5000	reasonable delay	No	Yes
LA	Louisiana	2006-01-01	0.5216	reasonable delay	Yes	Yes
NJ	New Jersey	2006-01-01	0.4167	reasonable delay	No	Yes
MN	Minnesota	2006-01-01	0.0833	reasonable delay	No	No
PR	Puerto Rico	2006-01-05	0.3549	1 day	No	No
ME	Maine	2006-01-31	0.3464	reasonable delay	No	Yes
OH	Ohio	2006-02-17	0.2500	reasonable delay	No	Yes
MT	Montana	2006-03-01	0.2500	reasonable delay	No	Yes
WI	Wisconsin	2006-03-31	0.1667	not specified	No	No
PA	Pennsylvania	2006-06-20	0.2500	reasonable delay	No	Yes
IL	Illinois	2006-06-27	0.3333	reasonable delay	No	Yes
ID	Idaho	2006-07-01	0.2081	reasonable delay	No	No
IN	Indiana	2006-07-01	0.3968	reasonable delay	No	Yes
NE	Nebraska	2006-07-14	0.2500	reasonable delay	No	Yes
CO	Colorado	2006-09-01	0.3333	reasonable delay	No	Yes
AZ	Arizona	2006-12-31	0.7867	45 days	No	Yes
HI	Hawaii	2007-01-01	0.5131	reasonable delay	No	Yes
KS	Kansas	2007-01-01	0.2500	reasonable delay	No	Yes
NH	New Hampshire	2007-01-01	0.4167	not specified	Yes	Yes
UT	Utah	2007-01-01	0.3918	reasonable delay	No	Yes
DC	District of Columbia	2007-07-01	0.6667	reasonable delay	Yes	Yes
WY	Wyoming	2007-07-01	0.5000	reasonable delay	No	Yes
MI	Michigan	2007-07-02	0.4167	reasonable delay	No	Yes
OR	Oregon	2007-10-01	0.3333	reasonable delay	No	Yes
MA	Massachusetts	2007-10-31	0.4167	reasonable delay	No	Yes
MD	Maryland	2008-01-01	0.5000	not specified	Yes	Yes
WV	West Virginia	2008-06-06	0.2500	reasonable delay	No	Yes
VA	Virginia	2008-07-01	0.3968	reasonable delay	No	Yes
IA	Iowa	2008-07-01	0.3333	reasonable delay	No	Yes
OK	Oklahoma	2008-11-01	0.4802	reasonable delay	No	Yes
TX	Texas	2009-04-01	0.4865	reasonable delay	No	Yes
AK	Alaska	2009-07-01	0.3833	reasonable delay	Yes	No
SC	South Carolina	2009-07-01	0.5018	reasonable delay	Yes	Yes
MO	Missouri	2009-08-28	0.4802	reasonable delay	No	Yes
MS	Mississippi	2011-07-01	0.4167	reasonable delay	No	Yes
VT	Vermont	2012-05-08	0.3333	reasonable delay	No	Yes
FL	Florida	2014-07-01	0.8283	30 days	No	Yes
KY	Kentucky	2014-07-15	0.0833	reasonable delay	No	No
NM	New Mexico	2017-06-16	0.4167	reasonable delay	No	Yes
AL	Alabama	2018-06-01	0.4583	45 days	No	Yes
SD	South Dakota	2018-07-01	0.5000	not specified	No	Yes
RI	Rhode Island	2023-06-27	0.2500	reasonable delay	No	No

Notes. The strictness index averages the six normalised legal dimensions described above. Dates follow the enforcement-date schedule merged into the firm-year panel for the DBN timing design.

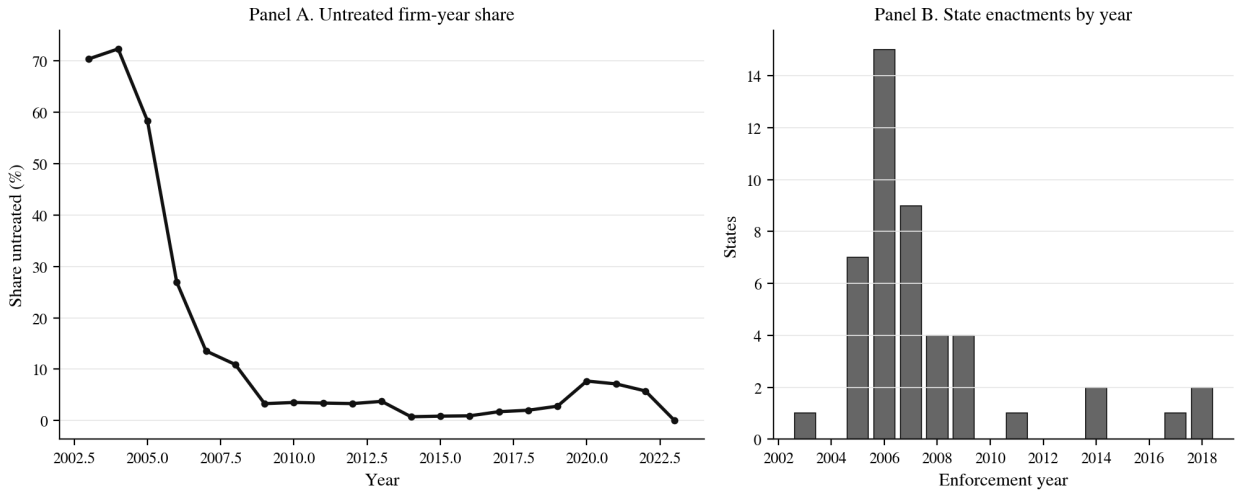

FIGURE A5. DBN treatment timing and untreated firm-year share

Notes. Panel A plots the share of firm-years that remain untreated in each year. Panel B plots state DBN law enactments by year. The peak enactment year is 2005, with 15 new laws. The untreated share falls from 58.4% in 2005 to 0.9% in 2015 and 0.0% in 2023.

A.5 LLM Prompt Text

This subsection reproduces the prompts used for the three LLM-based tasks. All prompts request JSON-formatted responses.

AI use-type classification prompt.. Used for 9,322 firm-years with non-zero AI activity.

> You are an expert analyst classifying how publicly listed US firms use artificial intelligence based on their annual business description from SEC 10-K filings.
>
> Read the business description below and classify the firm's AI activities. A firm may fall into multiple categories.
>
> CATEGORIES:

- DEV: The firm develops AI models, algorithms, or AI-powered products as a core business or R&D activity.
- ADOPT_3P: The firm adopts or integrates AI solutions developed by third parties.
- AI_CUSTOMER: The firm delivers AI-powered features to its customers or end users.
- AI_INTERNAL: The firm uses AI to improve its own internal operations.
- MENTION: AI or related terms appear but without substantive description of actual use or development.
- NONE: No meaningful AI-related content.

> INSTRUCTIONS: Assign ALL categories that apply. If the description mentions AI only in generic forward-looking statements, classify as MENTION. Focus on CURRENT activities, not aspirational plans. Be conservative: if uncertain between DEV and ADOPT_3P, look for whether the firm describes building vs. using.
>
> Respond ONLY with a JSON object: {"categories": ["CODE1", "CODE2"], "confidence": "high" | "medium" | "low", "evidence_snippet": "One sentence from the text", "ai_intensity": 0–10}

Causal-statement extraction prompt.. Used for 500 firm-years with non-zero AI activity and non-null risk-factor disclosures.

> You are an expert analyst examining risk factor disclosures from SEC 10-K filings.
>
> Read the risk disclosure below and identify any statements that explicitly link artificial intelligence, machine learning, data analytics, or advanced technology adoption to data breach risk, data security risk, or cybersecurity vulnerability.
>
> WHAT TO LOOK FOR: Statements where the firm says its use of AI or advanced technology creates, increases, or exposes it to data security risks; statements where AI systems are described as vulnerable to data breaches or cyberattacks; statements about third-party AI vendors creating data security exposure; statements about data requirements of AI increasing the firm's attack surface.
>
> WHAT TO EXCLUDE: Generic cybersecurity risk statements that do not mention AI or technology adoption as a cause; statements about AI being used to DEFEND against breaches (classify separately); boilerplate language about general technology risks.
>
> Respond ONLY with a JSON object: {"has_causal_link": true | false, "statements": [{"text": "Exact quote (max 50 words)", "direction": "AI_CAUSES_RISK" | "AI_MITIGATES_RISK" | "AMBIGUOUS", "specificity": "firm_specific" | "generic_boilerplate"}], "overall_tone": "defensive" | "proactive" | "neutral" | "no_relevant_content", "ai_risk_prominence": 0–10}

Validation scoring prompts.. Two validation prompts were used on a stratified sample of 1,000 firm-years. The AI-activity prompt asks for a 0–10 intensity rating of AI activity in the business description. The breach-risk prompt asks for a 0–10 prominence rating of data breach or data security risk in the risk disclosure. Both request JSON responses with a brief justification.

A.6 LLM and Pipeline Provenance

Component or task	Source layer	Sample	Role
Disclosure analysis main measures	Term-set matching	84,012 / 51,771	Lead invention, adoption, and breach-risk measures
Filing-level AI-use classification (full sample)	Filing-level LLM classification	9,322	Primary multi-category mechanism layer
Use-type validation benchmark	Validation-pass reclassification	500	Benchmark for filing-level classification
Direct-text extraction and statement reconciliation	Direct-text extraction passes	375	Appendix comparison and direct textual evidence in Section 6
Scalar validation scoring	Scalar validation prompts	1,000	AI-intensity and breach-prominence validation
Paragraph classification	Paragraph-level LLM classification	full panel	High-precision mechanism and appendix validation
DBN analysis outcomes	Term-set matching	83,392	Regulatory feedback analysis

TABLE A7. Measurement and LLM provenance

The filing-level AI-use classifier over 9,322 firm-years was run on 16 March 2026 with `claude-haiku-4-5` and `max_tokens = 300`. The overlapping use-type benchmark, the direct-text extraction pass, and the scalar validation pass were rerun on 18 March 2026 with `gpt-5.4-mini` through the Responses API; the respective output caps were 300, 500, and 200 tokens. The paragraph calibration sample and full paragraph classification run were also executed on 18 March 2026 with `gpt-5.4-mini`, `max_tokens = 220`, through the same API. The model split reflects pipeline timing rather than model shopping. The large full-sample filing-level pass pre-dated the later Responses API reruns, which were then standardised on `gpt-5.4-mini`.

The API calls set model, instructions, input, and `max_output_tokens` only; temperature was not set. Responses were required to return valid JSON. Post-processing stripped markdown code fences before JSON parsing, and responses that still could not be parsed were flagged as errors. For the filing-level classifier, deterministic parsing recovery yielded 9,321 usable rows out of 9,322. For the direct-text evidence, the main text reports the conservative reconciliation pass in Table A10, which retains only explicit AI-to-breach links after removing over-inclusive initial matches. Paragraph labels were then aggregated to firm-year indicators and shares by counting excerpt-level adoption, invention, and customer-facing units.

Exact agreement on the full category set in the benchmark sample is 67.3 per cent,

and category-level agreement for the substantive labels ranges from 88 to 95 per cent. Appendix A therefore documents why the paper uses the broad filing-level measures as the lead regressors and the narrower LLM layers as supporting mechanism and validation tools.

B. Additional Empirical Design and Identification Details

This section keeps only the supplementary DBN timing check that adds information beyond the main text. The broader lead-lag, realised-breach, and DBN power diagnostics are reported in Appendix C.2 and Appendix C.3.

B.1 Supplementary Event-Study Timing Checks

Table A8 reports the DBN event-study timing check without financial controls. The purpose is narrow. It asks whether the adoption-side negative path is entirely an artefact of the preferred control set. It is not. In the no-controls design, the invention path turns negative on impact and the adoption path becomes increasingly negative at longer horizons, consistent with the main-text ranking.

TABLE A8. Event-study timing check without financial controls

	AI invention measure	AI adoption measure
Event time −10	0.162	0.185
	(0.111)	(0.190)
Event time −9	0.107	0.151
	(0.100)	(0.173)
Event time −8	0.075	0.079
	(0.088)	(0.156)
Event time −7	0.045	0.070
	(0.076)	(0.133)
Event time −6	0.010	0.079
	(0.065)	(0.111)
Event time −5	0.033	0.072
	(0.057)	(0.092)
Event time −4	-0.005	0.061
	(0.047)	(0.069)
Event time −3	-0.025	0.037
	(0.036)	(0.048)
Event time −2	-0.007	0.019
	(0.025)	(0.026)
Event time −1	NA	NA
	(NA)	(NA)
Event time 0	-0.082*	-0.034
	(0.049)	(0.030)
Event time 1	-0.070*	-0.043
	(0.038)	(0.046)
Event time 2	-0.077*	-0.088
	(0.046)	(0.066)
Event time 3	-0.027	-0.158*
	(0.056)	(0.086)
Event time 4	-0.053	-0.158
	(0.066)	(0.105)
Event time 5	-0.053	-0.189
	(0.075)	(0.130)
Event time 6	-0.087	-0.234
	(0.086)	(0.155)
Event time 7	-0.115	-0.250
	(0.099)	(0.184)
Event time 8	-0.134	-0.273
	(0.108)	(0.198)
Event time 9	-0.102	-0.288
	(0.120)	(0.224)
Event time 10	-0.110	-0.344
	(0.131)	(0.247)
Financial controls	No	No
N	57,917	57,917

Notes. Statistical significance levels are * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

This check is supportive rather than decisive. It shows that the negative adoption path is not an artefact of one specific control set, but it does not overturn the limited-support caution that remains central to the DBN design.

C. Additional Results and Diagnostics

This section collects the supporting results that sit behind the main-text findings. Its purpose is to show which parts of the evidence sharpen interpretation, which parts qualify identification, and which parts should remain secondary.

C.1 Topic Specificity, Alternative Measures, and Statement Reconciliation

Table A9 collects the appendix diagnostics that matter most for construct validity. The preferred horse race keeps the adoption coefficient positive while the invention coefficient remains near zero in the lead filing-level specification. Removing vocabulary shared with the invention measure still leaves adoption positive at 0.0017. Adding the explicit non-AI digitisation placebo leaves the wedge intact, even though the placebo itself also loads positively and with a somewhat larger coefficient. That magnitude is expected because the placebo spans a broader set of digital-transformation language with more textual mass than the narrower adoption dictionary. Patent-based measures also load positively, which is consistent with the invention measure tracking real AI-related technical activity rather than only filing style.

Figure A6 gives the clearest visual summary of the topic-specificity test. Table A10 complements it with the remaining alternative-specification checks and the direct-text extraction comparison. In the raw baseline, 11 positive topics survive the BH correction. Once filing-length controls are added, the count falls to 2, and Topic 45 remains the top positive loading throughout, including when the invention measure is entered jointly and when the sample is restricted to the post-2010 period. Panel A of Table A10 reports the remaining alternative-specification checks. Panel C compares the initial direct-text extraction pass with the more conservative reconciliation pass reported in the main text. The directional imbalance is overwhelming in both runs.

TABLE A9. Measurement-cleanup and external-check diagnostics

Result	Estimate	Std. error	N
Panel A. Adoption-versus-invention wedge after measure cleanup			
Preferred horse race: AI invention measure	0.0000	0.0005	51,698
Preferred horse race: AI adoption measure	0.0021***	0.0004	51,698
Zero-overlap horse race: AI invention measure	0.0006	0.0004	51,698
Zero-overlap horse race: AI adoption measure	0.0017***	0.0004	51,698
Non-AI digitisation placebo horse race: AI invention measure	-0.0002	0.0004	51,746
Non-AI digitisation placebo horse race: AI adoption measure	0.0016***	0.0004	51,746
Non-AI digitisation placebo horse race: Digitisation placebo	0.0021***	0.0005	51,746
Panel B. Patent-side external checks			
Patent log measure	0.0021***	0.0006	18,103
Patent raw count	0.0000***	0.0000	18,103

Notes. Panel A shows that the filing-level adoption coefficient survives the appendix cleanup exercises that matter most for construct validity, including the explicit non-AI digitisation placebo. Panel B reports the patent-side external checks used to validate the invention margin. The patent-matched sample in Panel B adds two observations relative to the 18,101-observation full-controls disclosure sample. Significance levels are * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

TABLE A10. Topic specificity, alternative specifications, and direct-text comparison

Result	Estimate	Std. error	N
Panel A. Alternative measures and specifications			
Dictionary breach proxy, AI invention measure	0.0207	0.0129	27,630
Dictionary breach proxy, AI adoption measure	0.0272***	0.0102	27,630
Industry-year FE, AI invention measure	0.0007*	0.0004	18,101
Industry-year FE, AI adoption measure	0.0014**	0.0006	18,101
Winsorised+trimmed, AI invention measure	0.0004	0.0003	17,367
Winsorised+trimmed, AI adoption measure	0.0011**	0.0005	17,367
Panel B. Topic-specificity gate			
Full sample baseline	0.0025	BH+ topics: 11	Rank: 1
Full sample + length controls	0.0021	BH+ topics: 2	Rank: 1
Length controls + AI invention measure	0.0021	BH+ topics: 2	Rank: 1
2010-2023 + length controls + AI invention measure	0.0020	BH+ topics: 3	Rank: 1
Panel C. Direct-text extraction comparison			
Initial extraction pass	Links: 104 / 375	Raise: 237	Reduce: 4
Conservative reconciliation pass	Links: 45 / 375	Raise: 101	Reduce: 2

Notes. Panel A collects the alternative-specification checks that are not reported elsewhere. Panel B shows how filing-length controls narrow the number of positive BH-significant topics while Topic 45 remains the top positive loading. Panel C compares the initial statement-extraction run with the more conservative extraction reported in the main text.

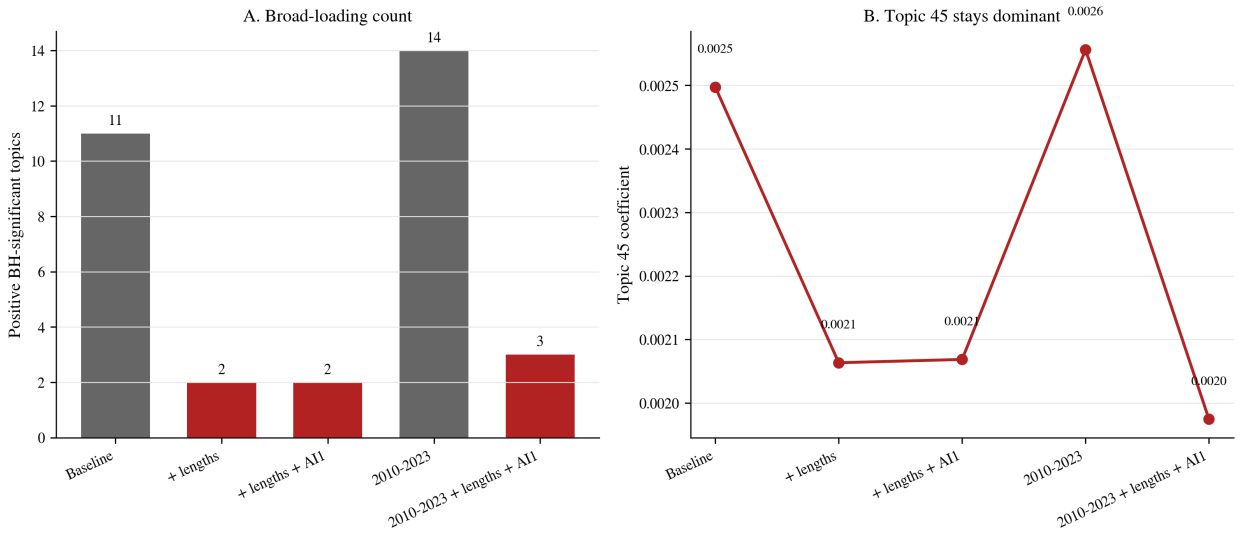

FIGURE A6. Topic-specificity gate

Notes. Panel A plots the number of positive BH-significant topics under five specifications. The count falls from 11 to 2 once filing-length controls are added. Panel B plots the Topic 45 coefficient, which remains the top positive loading throughout.

The paragraph-based specificity exercise is more supportive than decisive. In the full-sample baseline, Topic 45 has coefficient 0.000863 ($p = 0.0099$; BH-adjusted $p = 0.0355$), but 13 positive topics survive the multiple-testing correction. With length controls, the pattern narrows to 2 positive BH-significant topics, although Topic 45 no longer clears the BH threshold. The paper therefore uses the paragraph-specificity exercise only as a supporting diagnostic.

Table A11 addresses the concern that the disclosure result is driven by a small number of high-AI outliers. Dropping the upper tail, winsorising it, or switching to the extensive margin leaves the no-controls adoption effect positive and statistically strong.

TABLE A11. Tail robustness of the main adoption result

Specification	Estimate	Std. error	p-value	N
Baseline no-controls TWFE	0.002332	0.000532	< 0.001	50,867
Drop top 1% of AI adoption	0.002376	0.000335	< 0.001	50,339
Drop top 0.5% of AI adoption	0.002540	0.000388	< 0.001	50,599
Winsorise at 99th percentile	0.002818	0.000414	< 0.001	50,867
Extensive margin ($AI2 > 0$)	0.003744	0.000734	< 0.001	50,867

Notes: Results are from the tail-robustness diagnostics. The extensive-margin specification replaces the continuous adoption intensity with a binary indicator for any adoption language.

C.2 Timing and Realised-Breach Boundary Checks

Table A12 documents the supporting timing and boundary exercises. Lagged AI adoption predicts current breach-risk attention with a joint p-value of 0.060; the reverse direction is not jointly significant for either margin. Realised-breach regressions across LPM, Poisson, logit, and probit specifications produce small and imprecisely estimated coefficients, consistent with the severe underpowering from a 0.43 per cent breach-positive rate.

TABLE A12. Lead-lag timing and realised-breach boundary checks

Result	Estimate	Std. error	N
Panel A. Lead-lag timing checks			
Forward: breach-risk attention on lagged invention measure	L1 0.0005; L2 -0.0005*	Joint p = 0.134	14,415
Forward: breach-risk attention on lagged adoption measure	L1 0.0009**; L2 -0.0006	Joint p = 0.060	14,415
Reverse: invention measure on lagged breach-risk attention	L1 1.2775; L2 -0.4106	Joint p = 0.291	14,414
Reverse: adoption measure on lagged breach-risk attention	L1 0.9166; L2 0.1777	Joint p = 0.270	14,414
Panel B. Realised-breach boundary checks			
LPM - AI invention measure	-0.0005	0.0005	19,277
LPM - AI adoption measure	-0.0001	0.0006	19,277
Poisson - AI invention measure	-0.1198	0.1793	733
Poisson - AI adoption measure	0.0690	0.1828	733
Logit - AI invention measure	-0.1083	0.2149	19,277
Logit - AI adoption measure	-0.0120	0.2214	19,277
Probit - AI invention measure	-0.0386	0.1005	19,277
Probit - AI adoption measure	0.0035	0.1066	19,277

Notes. These models are useful as supporting evidence only. The adoption lead-lag is directionally supportive, and the realised-breach models are underpowered because breach events are rare in the matched panel.

Tables A13 and A14 report the detailed timing and realised-breach layouts.

TABLE A13. Lead-lag specifications

Model	Lag 1 term	Lag 1 est.	Lag 1 p	Lag 2 term	Lag 2 est.	Joint p
DB score on lagged AI invention	$AI1_{t-1}$	0.00054	0.1826	$AI1_{t-2}$	-0.00050	0.1337
DB score on lagged AI adoption	$AI2_{t-1}$	0.00095	0.0232	$AI2_{t-2}$	-0.00064	0.0598
AI invention on lagged DB score	DB_{t-1}	1.27746	0.1948	DB_{t-2}	-0.41056	0.2910
AI adoption on lagged DB score	DB_{t-1}	0.91656	0.1060	DB_{t-2}	0.17770	0.2695

Notes: All models include firm and year fixed effects and the lagged financial controls used in the main controlled specification.

TABLE A14. External realised-breach boundary checks

Model	AI adoption estimate	Std. error	*p*-value	*N*
DBC conservative LPM	-0.0003	0.0019	0.8625	16,784
DBC conservative LPM, no controls	0.0010	0.0016	0.5540	44,580
DBC forward 3-year LPM	-0.0046	0.0028	0.1009	14,031
DBC hack-only LPM	-0.0021	0.0012	0.0801	16,784
ITRC conservative LPM	0.0015	0.0013	0.2432	10,987

Notes: Results are summarised from the external-breach regression pass. These specifications do not overturn the disclosure result, but they do not validate it either. We therefore treat realised-breach evidence as a boundary check rather than as a main empirical pillar.

C.3 DBN Robustness and Geographic Concentration

Table A15 and Figure A7 document the stability and geographic structure of the DBN result. The sign remains negative under state-clustered standard errors, when early-adopter jurisdictions are excluded, and when headquarters movers are dropped. The jackknife mean across 49 iterations is −0.229 (SD = 0.033). California is the single most consequential state. Removing it shifts the estimate by 74 per cent of the baseline magnitude. The ex-California estimate remains directionally negative but materially weaker. Removing Texas or Massachusetts makes the estimate more negative. The adoption effect is concentrated in a limited set of influential states rather than uniformly distributed.

Table A16 adds the Roth-style power slopes behind the pre-trend discussion. The adoption design remains the cleaner causal margin. The pre-trend joint statistic is smaller, the *p*-value is far from conventional rejection thresholds, and the detectable slopes under 50% and 80% power are modest relative to the realised post-treatment movement.

TABLE A16. Roth-style pre-trend sensitivity

Outcome	Slope for 50% power	Slope for 80% power	Pre-trend joint stat.	Pre-trend *p*
AI invention	0.0531	0.0822	15.6122	0.0754
AI adoption	0.0583	0.0924	5.8846	0.7514

Notes: Output from the `pretrends` package based on the DBN event-study specification.

TABLE A15. DBN robustness and design diagnostics

Result	Estimate	Std. error	N
Panel A. Alternative DBN specifications for AI adoption			
Baseline TWFE adoption outcome	-0.230	E3 -0.299**	E5 -0.384*
State-clustered SE	-0.230	E3 -0.299*	E5 -0.384
Exclude early-adopter states	-0.060	E3 -0.095	E5 -0.216
Exclude HQ movers	-0.079	E3 -0.113	E5 -0.125
Panel B. Design diagnostics			
AI invention pretrend p-value	0.075		
AI adoption pretrend p-value	0.751		
Jackknife mean post average	-0.229	SD 0.033	States 49
Post average without California	-0.060	Shift 74.0%	
States with >50% shift	CA		

Notes. Panel A reports the adoption-side DBN sensitivity variants that remain useful for interpretation. Panel B records the pretrend diagnostics and the leave-one-state-out jackknife summary, which shows that the signal remains directionally negative but is materially weaker without California.

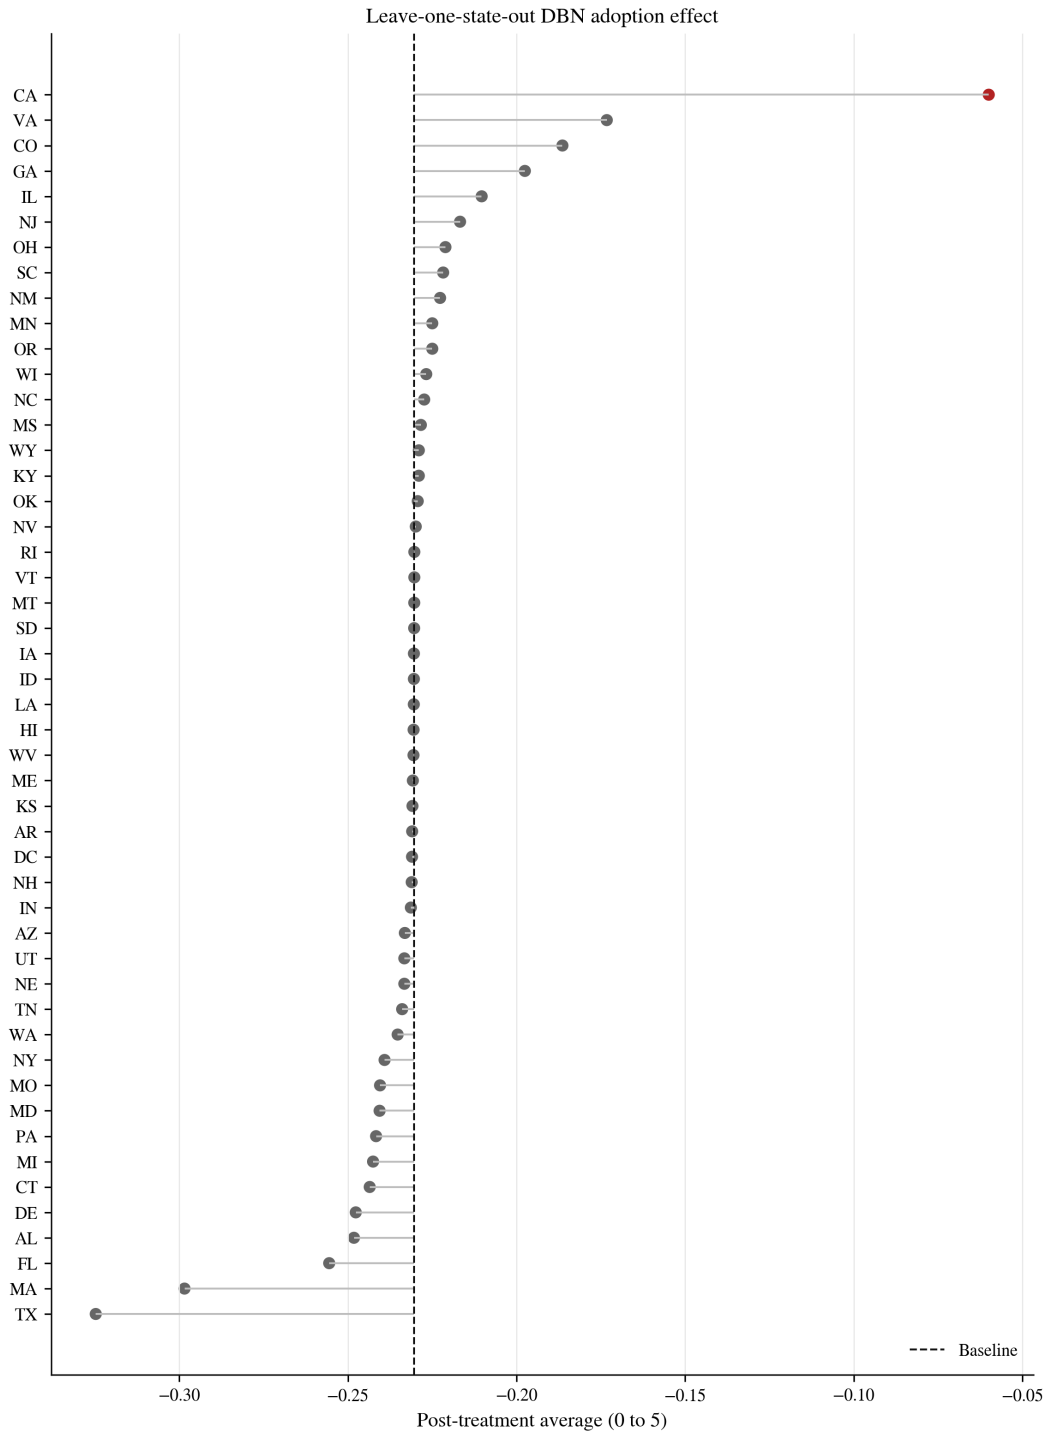

FIGURE A7. Leave-one-state-out DBN adoption effect

Notes. Each row shows the TWFE average post-treatment effect (event times 0 through 5) for AI adoption when one state is excluded. The dashed vertical line marks the baseline estimate of −0.230. California produces the largest attenuation. Removing Texas or Massachusetts makes the estimate more negative.

C.4 Paragraph Measure and Heterogeneity

Table A18 documents why the paragraph measure does not replace the preferred filing-level measure as the lead regressor. In the full-sample horse race, paragraph adoption is positive (0.0008, $p = 0.014$) while paragraph invention is not statistically distinguishable from zero. In the post-2012 subsample, the invention coefficient rises to 0.0005 ($p = 0.039$), compressing the wedge. One plausible interpretation is that later AI invention itself becomes more data-intensive, so the liability distinction between invention and adoption narrows. The paragraph Mundlak decomposition shows a dominant between-firm adoption component (0.338) and a positive but smaller within-firm component (0.088). In the horse-race decomposition, the between-firm invention component is negative (–0.219, $p < 0.05$), but that sign reversal should be read cautiously because the paragraph invention series is extremely sparse and therefore identifies a much narrower set of firms than the filing-level measure.

The heterogeneity estimates in Panel C and Figure A8 show that the positive adoption coefficient appears across most subgroups, with the clearest precision in technology firms (0.0013), small firms (0.0015), and both low- and high-baseline cybersecurity groups (0.0015 and 0.0021). The heterogeneity evidence is useful for showing that the main association is not confined to a single sector, but it is not strong enough to support new primary claims about subgroup treatment effects.

Table A17 restores the full paragraph-specification ladder that sits behind the compressed horse-race rows in Table A18. The broader no-controls adoption result survives, but the stricter measure is more vulnerable to over-control and later-period compression than the lead filing-level series.

TABLE A17. Paragraph-based measure for alternative disclosure results

Specification	Estimate	Std. error	p-value	N
Full sample, no controls, adoption	0.000863	0.000335	0.0099	51,771
Full sample, log assets only, adoption	0.000240	0.000316	0.4464	46,985
Full sample, all controls, adoption	0.000126	0.000536	0.8143	18,103
2012–2023, no controls, adoption	0.000949	0.000429	0.0272	34,836
Full sample horse race, adoption	0.000820	0.000333	0.0138	51,771
Full sample horse race, invention	0.000324	0.000233	0.1646	51,771
2012–2023 horse race, adoption	0.000883	0.000429	0.0396	34,836
2012–2023 horse race, invention	0.000510	0.000247	0.0391	34,836

Notes. Coefficients are reported in raw disclosure units. The stricter anchor set preserves the positive adoption sign in the broad full-sample regressions, but it does not preserve the invention-null wedge as cleanly in later-period horse races.

TABLE A18. Paragraph-measure and heterogeneity checks

Result	Estimate	Std. error	N
Panel A. Paragraph-measure horse races			
Full sample horse race: Adoption	0.0008**	0.0003	51,723
Full sample horse race: Invention	0.0003	0.0002	51,723
2012-2023 horse race: Adoption	0.0009**	0.0004	34,251
2012-2023 horse race: Invention	0.0005**	0.0002	34,251
Panel B. Paragraph-measure Mundlak support			
Paragraph Mundlak simple: Within adoption	0.088***	0.020	51,770
Paragraph Mundlak simple: Between adoption	0.338***	0.053	51,770
Paragraph Mundlak full controls: Within adoption	0.028	0.024	18,102
Paragraph Mundlak full controls: Between adoption	0.394***	0.062	18,102
Paragraph Mundlak horse race: Within invention	0.092*	0.047	51,770
Paragraph Mundlak horse race: Within adoption	0.084***	0.020	51,770
Paragraph Mundlak horse race: Between invention	-0.219**	0.094	51,770
Paragraph Mundlak horse race: Between adoption	0.349***	0.055	51,770
Panel C. Heterogeneity summary			
Tech	0.0013**	0.0007	3,762
Finance	0.0017	0.0020	346
Healthcare	-0.0002	0.0025	153
Manufacturing	0.0013	0.0009	12,570
Other	0.0038	0.0027	1,270
Small firms	0.0015**	0.0006	10,708
Large firms	0.0019*	0.0010	7,393
Low baseline cybersecurity	0.0015**	0.0006	15,308
High baseline cybersecurity	0.0021**	0.0010	2,793

Notes. The paragraph measure is reported as a narrower supporting layer rather than the lead disclosure measure. The heterogeneity rows are descriptive second-order checks rather than a central identification margin.

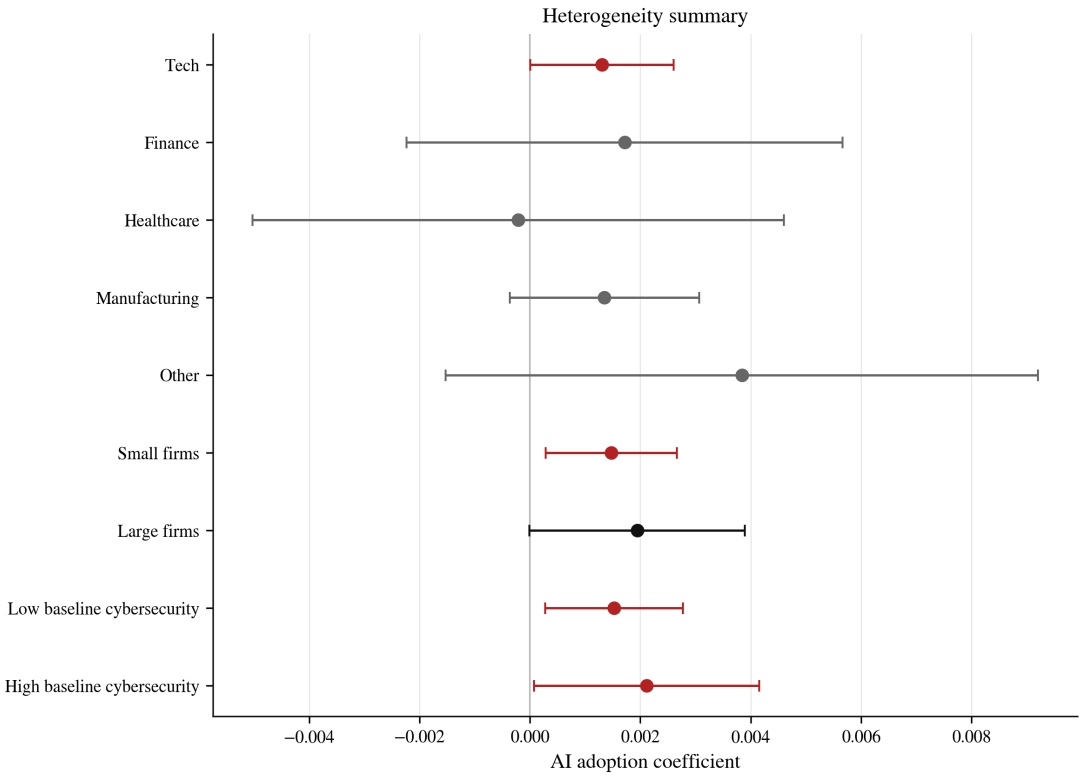

FIGURE A8. Heterogeneity summary

Notes. Each row plots the AI adoption coefficient from a separate subsample regression. Red markers indicate significance at the 5 per cent level; grey markers indicate insignificance. The clearest positive estimates are for technology firms, small firms, and both low- and high-baseline cybersecurity groups.

The appendix results do not change the paper's main reading. They narrow it and make it more defensible. The benchmark adoption coefficient becomes more topic-specific once filing-level verbosity is absorbed. The paragraph and direct-text layers reinforce the interpretation that the relevant margin is operational deployment. The DBN evidence remains negative on the adoption side but is concentrated in a limited set of influential states. The heterogeneity evidence confirms that the association is not confined to a single sector without supporting sharper subsample claims.