INTERNATIONAL MONETARY FUND

Jurisdictional Capital and AI Regulation: Evidence from the EU AI Act

Yi Chen, Zhe Wang, Jing Zhou

WP/26/180

IMF Working Papers **describe research in progress by the author(s) and are published to elicit comments and to encourage debate.**
The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

2026
AUG

WORKING PAPER

© 2026 International Monetary Fund WP/26/180

IMF Working Paper
Western Hemisphere Department

Jurisdictional Capital and AI Regulation: Evidence from the EU AI Act

Prepared by Yi Chen, Zhe Wang, Jing Zhou*

Authorized for distribution by Nigel Chalk
August 2026

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: We study how AI regulation affects firm valuation using the EU Artificial Intelligence Act, the world's first comprehensive AI framework. In an event study around the April 2021 proposal, we find firms combining deeper EU presence with faster AI hiring earned higher announcement returns, suggesting markets value "jurisdictional capital"—experience in the EU regulatory environment helps firms navigate the AI regulation. The effect is stronger for high-risk AI, for firms with stable and concentrated EU presence, or prior compliance experience, unexplained by size, foreign exposure, or lobbying. EU-embedded, AI-expanding firms increase within-firm EU revenue share when peers are less embedded.

JEL Classification Numbers:	G14, G18, K23, O33
Keywords:	AI regulation, event study, EU AI Act, jurisdictional capital
Author's E-Mail Address:	yc2535@cornell.edu, zhewang192@gmail.com, jzhou@imf.org

* We thank our corresponding institutions for financial support. We thank JaeBin Ahn, Andrew Caplin, Nigel Chalk, Mariarosaria Comunale, Will Cong, Mai Chi Dao, Adriano Fernandes, Niels-Jakob Hansen, Peter Iliev, Ioannis Lianos, Ran Shorrer, Sneha Thube, Alice Tianbo Zhang, and the Bank of Spain for helpful comments. All errors are our own. The views expressed here are those of the author and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

WORKING PAPERS

Jurisdictional Capital and AI Regulation: Evidence from the EU AI Act

Prepared by Yi Chen, Zhe Wang, and Jing Zhou

1 Introduction

Rapid development of Artificial Intelligence (AI) has intensified the case for regulations that address its safety, accountability, and risks. AI regulation can increase legal certainty, reduce market fragmentation, and support trust in AI systems, meanwhile, firms differ in their ability to navigate and implement new regulatory requirements. This paper asks: How does AI regulation affect firm valuation, and which firms are better positioned to adapt?

We provide a first look at these questions using the proposal announcement of the EU Artificial Intelligence Act (EU AI Act), the world's first comprehensive AI legislative framework. Our key finding is that, after the proposal was announced, Firms combining deeper EU market presence with faster pre-event AI hiring growth experienced relatively higher announcement returns. We interpret this effect as market pricing *jurisdictional capital*: the accumulated experience, routines, and institutional knowledge that firms build through sustained participation in a regulatory jurisdiction, which can help firms to navigate the new AI regulations.

The EU AI Act offers a natural setting for studying how AI regulation is perceived by private market. Proposed by the European Commission on April 21, 2021 and finalized in December 2023, the Act establishes risk-based compliance obligations that vary by AI use-case risk category, imposing requirements such as conformity assessments, algorithmic audits, and mandatory human oversight. In addition, it has extraterritorial reach: any firm using or supplying AI products on the EU market must comply, regardless of where it is headquartered. U.S. firms therefore vary along two dimensions that jointly determine their regulatory exposure: reliance on AI technology (regulatory scope) and depth of EU market presence (regulatory jurisdiction). EU market presence could represent two aspects: it determines the Act's jurisdictional reach over the firm, and it proxies for accumulated jurisdictional capital under EU regulation.

Our identification strategy uses predetermined exposure measures constructed from pre-2021 data. We proxy for EU market presence by the average share of revenue earned in the EU, and for AI capability build-up by the growth in AI-related hiring intensity, both computed over 2015–2019 (the pandemic year 2020 is excluded). The valuation impact is measured by cumulative abnormal returns (CARs)—stock returns net of those predicted by aggregate market movements—over a short window around the proposal release date. The key variable is the interaction of EU revenue share and AI hiring growth, which captures the differential market response for firms simultaneously exposed to the Act's substantive scope (AI activity) and firm's jurisdictional capital (EU market presence).

Our key finding, a positive coefficient on the interaction term, indicates that jurisdic-

tional capital helps firms navigating the Act's requirements. Quantitatively, a one-standard-deviation increase in AI hiring growth raises the marginal effect of EU market presence on the three-day CAR by approximately 40 basis points. Translating to nominal values, Within AI-intensive firms, higher EU jurisdictional capital is valued about $86 billion higher in shareholder value at announcement.

We interpret this as evidence of jurisdictional capital from EU commitment: firms with deep, sustained presence in EU have built compliance infrastructure and interpretive knowledge that could help in fulfilling the Act's requirements. This aligns with market commentary as well, for instance, Fitch notes that firms with mature AI governance can absorb the AI Act's requirements with limited disruption.[1]

We next show that the jurisdictional capital is not driven by firm size, lobbying, or generic international diversification. High EU revenue share could capture scale or global reach rather than jurisdictional capital, and firms with deep EU roots may also have stronger political access to shape the Act in their favor. The evidence does not support these alternatives. Replacing EU revenue share with absolute EU revenue, firm size, or total foreign revenue share does not generate a comparable AI interaction. Physical capital intensity, if anything, increases the AI compliance cost, suggesting that operational scale is positively associated with regulation requirements. Direct measures of political access at the EU and member state levels do not absorb the EU×AI coefficient and do not independently predict abnormal returns. These results distinguish the AI Act from the EU General Data Protection Regulation (GDPR) setting, where the regulatory benefit is often interpreted as size based. Here, the relevant determinant is EU market embeddedness.

We then examine which firm characteristics strengthen the jurisdictional capital. The effect is stronger for firms expanding hiring in high-risk AI jobs, where the Act imposes the most strict obligations. It is also stronger among firms with stable EU revenue shares, geographically concentrated EU revenue, and prior EU related compliance experience. These patterns suggest that the jurisdictional capital comes from durable operating depth in the EU regulatory environment, not from broad or passive foreign exposure.

Beyond valuation effects, we examine whether jurisdictional capital translates into real economic outcomes. To this end, we examine the within-firm geographic composition of revenue. We find that, among firms with high AI hiring growth and high prior EU exposure, weaker EU-embeddedness among AI exposed peers predicts subsequent revenue reallocation toward the EU within the firm. Moving from a firm whose peers have median EU-embeddedness to one whose peers have no EU-embeddedness is associated with about a 21 percent larger expansion of EU revenue relative to non-EU revenue. For the median firm

[1] Fitch (05 Nov, 2025).

within the high EU and high AI sample, this corresponds to a 2.8 percentage point increase in EU revenue share. These patterns suggest that priced jurisdictional capital is reflected not only in valuation, but also in subsequent real reallocation.

Our findings show that the stock market prices jurisdictional capital as an intangible asset in reacting to the Act's proposal. Because stock-market reactions capture expected private cash flows and discount rates, our estimates provide a measure of how private markets evaluate the regulation. This private-value lens is complementary to the broader channels through which the Act operates, including consumer surplus, innovation, risk reduction, accountability, trust, legal certainty, and the gains from regulatory harmonization. Establishing this welfare-optimal scope of the Act would require modeling the externalities it targets and the surplus it generates, which is a distinct and worthwhile exercise that lies beyond an announcement event study of private valuations.

Related Literature. A large literature establishes that regulation can have different impact across firms (Stigler, 1971; Klapper, Laeven and Rajan, 2006), and that political and relationship capital is priced into firm value (Fisman, 2001; Faccio, 2006). We shift the locus of regulatory capital from political connections to jurisdictional commitment: the procedural knowledge and institutional fluency accumulated through sustained market participation, distinct from access to policymakers. Our results provide, in AI regulation, the first empirical evidence of jurisdiction capital through regulatory engagement.

The closest empirical precedent is the EU GDPR literature, which mostly documents a *scale-based* benefit: large firms are less negatively affected by the regulation and capture market share, while small players face more severe losses. The leading mechanisms emphasize fixed compliance costs and complementary advantages of large incumbents (e.g., brand strength, organizational capacity), which make compliance easier to absorb at scale (Johnson, Shriver and Goldberg, 2023; Peukert et al., 2022; Goldberg, Johnson and Shriver, 2024). We document a new mechanism under the EU AI Act. Here, the relevant determinant is jurisdictional commitment: EU market embeddedness predicts higher abnormal returns. This pattern is consistent with the Act imposing product- and use-case-specific requirements that scale with the number of high-risk systems and with the greater interpretive and implementation uncertainty of a new regulatory regime. The relevant determinant, therefore, is not scale, but EU-embedded operational capital—the ability to navigate and implement regulatory requirements within EU institutions.

Our findings also contribute to the literature on cross-jurisdictional regulation, but the mechanism we identify is distinct. Christensen, Hail and Leuz (2016) show that common EU securities rules have uneven effects because regulatory outcomes depend on country-level

prior conditions, implementation, and enforcement. We shift the locus of heterogeneity from regulatory institutions to regulated firms: under a common EU rule, firms differ in their accumulated capacity to operate within the regulator's jurisdiction. Zeume (2017) shows that unilateral anti-bribery regulation penalizes exposed firms and shifts business toward less constrained competitors. In contrast, we show that among firms exposed to the EU AI Act, deeper EU-embeddedness mitigates compliance costs and predicts subsequent reallocation toward the regulated market. Klein, Ludwig and Spengel (2022) show that EU digital tax proposals impose larger losses on firms with greater EU exposure, consistent with exposure increasing vulnerability to regulatory burden. We show that under AI regulation, the same jurisdictional footprint can play a second role: it is not only regulatory incidence, but also adaptation capability. The positive EU×AI interaction is consistent with the second role becoming relatively valuable.

Lastly, existing work on AI and firm value documents that AI investment is associated with higher growth and market valuations (Babina et al., 2024), that generative-AI workforce exposure generated positive returns after ChatGPT (Eisfeldt, Schubert and Zhang, 2023), and that technical versus ethical AI standards have opposite investment effects (Canayaz and Wang, 2026). We provide, to our knowledge, the first event study of how comprehensive AI regulation affects firm value according to firms' pre-existing AI and jurisdictional exposure. The April 2021 event predates the generative-AI boom initiated by ChatGPT, so our estimates should not be mechanically extrapolated to the current regulation of frontier models. The distinction is especially relevant as frontier models create new cyber and systemic-risk concerns: in June 2026, the European Systemic Risk Board warned that highly capable frontier models could increase the probability and severity of cyber incidents with systemic effects.[2]

The remainder of the paper is organized as follows. Section 2 provides institutional background on the EU AI Act. Section 3 describes the data. Section 4 presents the empirical design. Section 5 reports the baseline results. Section 6 and Section 7 test competing mechanisms and amplifiers, respectively. Section 8 explores the implication on within-firm EU market reallocation. Section 9 discusses implications for investors, regulators and practitioners. Section 10 concludes.

[2]European Systemic Risk Board, "Warning on systemic cyber risks stemming from frontier artificial intelligence models", ESRB/2026/3, 25 June 2026.

2 Institutional Background of EU AI Act

The EU AI Act is the world's first comprehensive legislative framework for artificial intelligence. We describe the features relevant to our identification strategy: the risk-based compliance architecture, the legislative timeline we exploit for event dates, and the Act's extraterritorial reach.

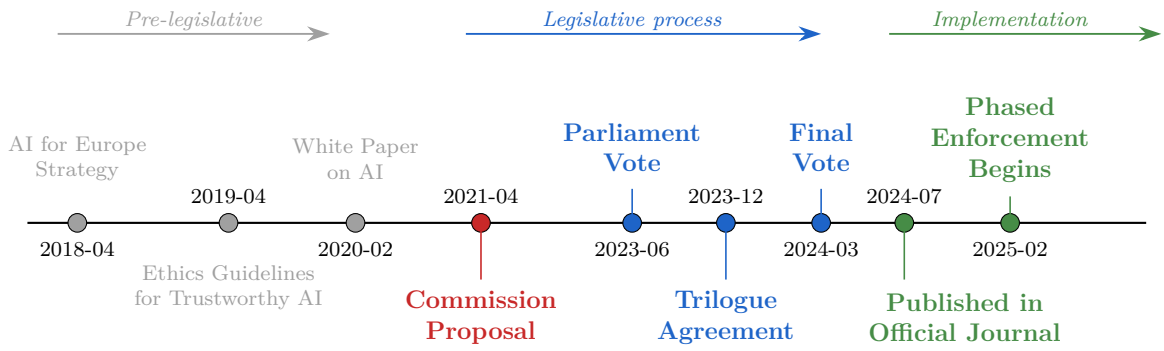


Figure 1: EU AI Act Timeline

Notes: The timeline shows the key milestones in the development of the EU AI Act. Gray markers indicate pre-legislative milestones that established the policy direction. The red marker denotes the European Commission's April 2021 proposal—our primary event date. Blue markers indicate subsequent legislative milestones used in the multi-event analysis. Green markers indicate adoption and the beginning of phased enforcement.

Risk-based compliance architecture. On April 21, 2021, the European Commission published its "Proposal for a Regulation laying down harmonised rules on artificial intelligence" (COM/2021/206).[3] The proposal classifies AI systems into four use-case risk tiers with correspondingly differentiated obligations: *unacceptable-risk* systems (AI use-cases such as manipulative AI, social scoring, "real time" remote biometric identification) are prohibited outright; *high-risk* systems (e.g., AI systems used in critical infrastructure, education, evaluation of creditworthiness) face mandatory conformity assessments, risk-management obligations, data-governance requirements, and human-oversight mechanisms;[4] *limited-risk* systems (e.g., chatbots) face transparency obligations; and *minimal-risk* systems face no additional requirements. Non-compliance penalties reach up to 7% of global annual turnover.

Legislative timeline. The Act's development from proposal to adoption spanned four milestones, which we use as event dates. The *April 21, 2021* Commission proposal introduced the risk-based architecture, the extraterritorial scope, and the enforcement regime. The *June*

[3] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206
[4] https://artificialintelligenceact.eu/article/6/

14, 2023 European Parliament vote expanded the Act's scope to cover foundation models and general-purpose AI—categories absent from the original proposal. The *December 9, 2023* trilogue agreement among the European Parliament, the Council presidency, and the European Commission locked in the final regulatory architecture. The *March 13, 2024* Parliament vote approved the Act, which was published in the Official Journal on July 12, 2024 with phased enforcement beginning in 2025. The April 2021 proposal is our primary event because it was the point at which the Act's cross-sectional incidence on firms was first revealed; subsequent milestones primarily refined an already-disclosed framework.

Extraterritorial reach. The Act applies to any providers or deployers that place AI systems on the EU market, regardless of where it is established. U.S. firms with EU revenue exposure are therefore directly subject to its requirements, which is the feature that makes our cross-sectional identification strategy feasible.[5]

3 Sample and Variable Construction

We combine six data sources covering the 771 US-listed firms in our main analysis sample. AI exposure is measured from Lightcast database of job postings that classifies each posting into AI-related categories, linked to Compustat by `gvkey`. EU revenue share is from FactSet's Geographic Revenue (GeoRev) database. Daily stock returns are from CRSP/Compustat Merged; Fama–French factors are from Kenneth French's data library; firm financial controls are from Compustat Annual. Appendix A describes data sources, the variable definitions, and the matching procedures in detail.[6]

Our two identifying variables are constructed from predetermined, pre-COVID data. *AI hiring growth* is the annualized change in firm i's AI hiring share between 2015 and 2019, where the AI hiring share in a given year is the share of AI-related postings in total postings (requiring at least five total postings in 2015 and 2019). We use AI hiring growth rather than the level because AI adoption was still at an early stage when the EU AI Act was first

[5] In contrast to the EU, comprehensive AI regulation elsewhere emerged later and more incrementally. U.S. regulation developed largely bottom-up, from early state laws to later federal frameworks like the Executive Order on AI (October 2023). Other major AI regulatory development, such as in the UK and Brazil, was introduced only in 2023 afterwards. None of these coincides with our April 2021 event window, leaving a clear gap around the announcement we study. In addition, the EU, UK, U.S., and Brazilian approaches differ significantly in design, notably, the EU AI Act's risk assessment approach. Brazil followed a similar path. The UK has not had any regulation framework definable as such to this day. The U.S. approach of the EO was not really designed to be a framework but rather as guidelines (Comunale and Manera (2024)). Given the difference across different jurisdictions, assessment of other AI regulations may require different empirical designs.

[6] Additional data used in robustness checks and extensions will be described in the corresponding sections.

proposed in 2021, and both the technology and firms' AI capabilities were evolving rapidly. Since the Act would become binding only several years later, hiring growth better captures the firm's forward looking AI capacity: the growth option that is likely to be exposed to the new regulatory requirements when the regulation is ultimately implemented (more on this in Section 5.4).

EU revenue share is the average share of firm revenue derived from the European Union in total revenue over 2015–2019. For each firm-year, the total revenue from EU region is the summation of disclosed revenue from each EU member country, excluding UK (due to Brexit in 2020). Firms with no disclosure in the window are dropped (we will turn to this extensive margin in Section B.8), and firms with partial disclosure contribute their mean over disclosed years.

One note is that the 2015-2019 period is used to construct the right-hand-side exposure variables (AI hiring growth and EU revenue share), not to measure the effects of the Act. The left-hand-side variable is the cumulative abnormal return in a narrow event window around the April 2021 proposal announcement, the date at which the regulation became public information and forward-looking investors could price its anticipated obligations (more in Section 4). Using pre-event, pre-COVID data for the exposure measures is a deliberate identification choice that avoids reverse causality and ensures the covariates are not themselves affected by the event.

Financial controls (log market capitalization, book-to-market, leverage, R&D intensity, profitability) are averaged over 2015–2019 to match the exposure measurement period; construction details are in Appendix A.2. The analysis sample comprises the 771 US-listed firms that satisfy six sequential filters—CRSP/Compustat match, disclosed EU revenue, sufficient job-posting coverage to compute AI growth, valid Fama–French CARs, non-missing controls, and a pre-event share price of at least \$5—with the attrition waterfall, matching procedure, and the rationale for retaining financial and utility firms documented in Appendix A.3.

Figure 2 plots the time-series evolution of the two exposure variables. The left panel shows that AI hiring was negligible through 2015—median AI share below 2%—before accelerating sharply: the median quadruples from 2% in 2015 to 8% in 2024, and the interquartile range widens as firms diverge in their adoption trajectories. The acceleration begins well before the April 2021 proposal, confirming that our 2015–2019 growth measure captures a predetermined trend rather than a response to the Act. The right panel shows that EU revenue share, by contrast, is more persistent: the median declines gradually from roughly 14% in 2005 to 9% in 2024, reflecting the secular growth of US domestic and Asian markets relative to the EU. This persistence is consistent with interpreting EU share as a measure of durable jurisdictional commitment rather than a transient revenue fluctuation.

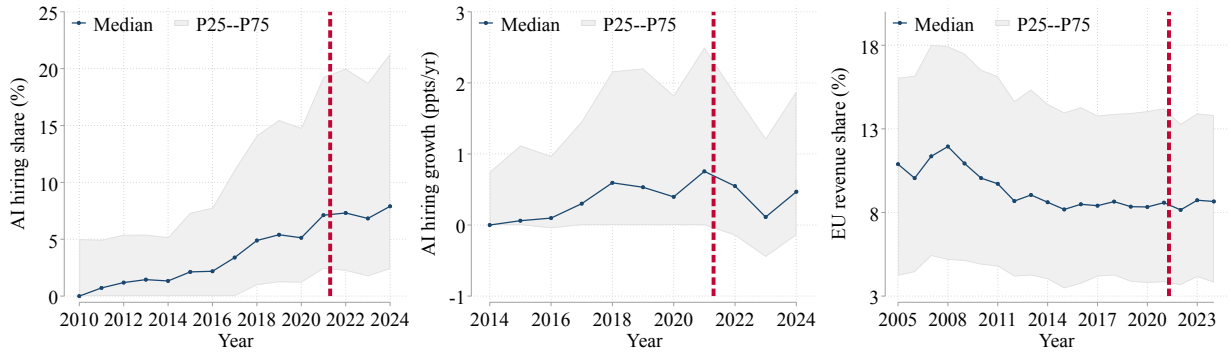


Figure 2: Time Series of Key Exposure Variables

Notes: Left and middle panel plot the median AI job posting share and AI job posting share annualized growth (blue line) with interquartile range (shaded) across the sample 2010–2024, respectively. Right panel plots the median EU revenue share (blue line) with interquartile range, 2005–2024. The dashed red line marks the EU AI Act proposal (April 2021).

Panel a) of Figure 3 displays the cross-sectional joint distribution of the two variables, colored by Fama–French 48 industry groups, while in panel b) and c), each point represents a sector average or a selected large technology firm, respectively. Two features stand out. First, the correlation between EU share and AI hiring growth is very limited ($\rho = -0.04$), so they capture distinct sources of regulatory exposure. Second, the distribution is concentrated in the lower-left quadrant (moderate EU share, low AI growth), with a long right tail of EU-intensive firms (chemicals, pharmaceuticals) and a long upper tail of AI-intensive firms (finance and business services). High-AI-growth firms such as NVIDIA stand in the upper-left, while high-EU-share firms such as Adobe anchor the right side. The orthogonality of the two exposure dimensions provides identifying variation for the interaction specification that follows. Summary statistics are reported in Table A3.

4 Empirical Design

4.1 Estimation Method

Our identification strategy relies on two elements. First, firms' ex ante exposure to the Act is proxied by pre-2021 EU revenue share and AI hiring growth. Second, a narrow event window around the April 2021 proposal isolates announcement-driven repricing. The baseline specification estimates a cross-sectional regression of cumulative abnormal returns on firms' EU revenue share and AI hiring growth.

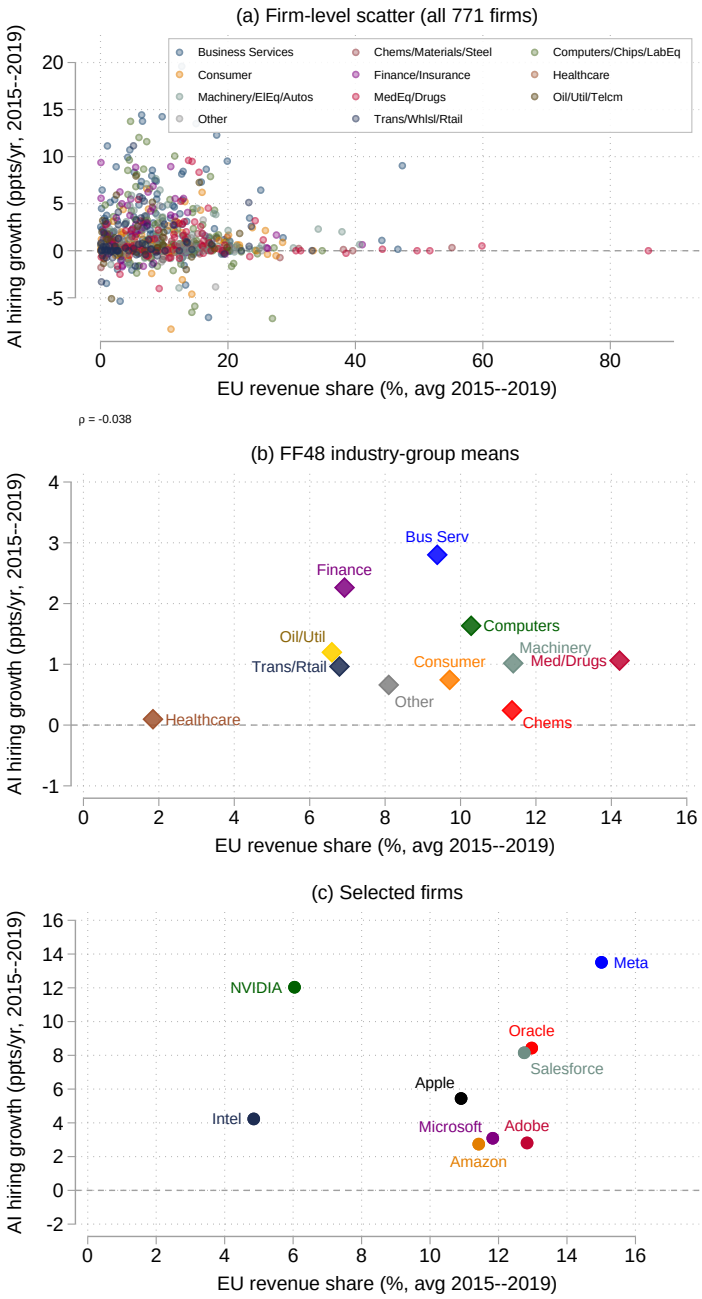


Figure 3: AI Hiring Growth vs. EU Revenue Exposure by FF48 Industry Group

Notes: In panel a, each point represents one of the 771 sample firms, plotted by EU revenue share (avg FY2015–2019, horizontal axis) and AI hiring growth (annualized change CY2015→2019, vertical axis). Colors indicate Fama–French 48 industry groups. The correlation between the two measures is $\rho = -0.038$. In panel b and c, each point represents a sector (labeled) or a selected large technology firm (labeled), respectively.

The key regression specification is as follows:

$$\text{CAR}_i[\tau_1, \tau_2] = \alpha + \beta_1(\text{EU} \times \text{AI})_i + \beta_2 \text{AI}_i + \beta_3 \text{EU}_i + \Gamma \mathbf{X}_i + \delta_s + \varepsilon_i, \tag{1}$$

where EU (i.e., EU revenue share) and AI (i.e., AI hiring growth) are standardized to mean zero and unit variance within the analysis sample, $\mathbf{X}_i$ denotes the financial controls described in Section 3, and δ_s are Fama-French-48 industry fixed effects.

The interaction term of EU and AI captures distinct dimensions of regulatory relevance. EU revenue share measures the historical importance of the EU market to the firm—how deeply committed the firm is to the EU jurisdiction. A high EU share signals that the firm has built its business around the EU market, investing in local operations and compliance infrastructure over time. AI hiring growth captures the firm's accumulation of AI-related growth options. Because the EU AI Act will take years to fully implement, the market should respond most to firms that are actively building AI capabilities for future deployment—not to firms with static AI operations that are already priced in. AI growth thus identifies firms whose future AI deployment options are most directly exposed to the new regulatory regime.

The key hypothesis is $\beta_1 > 0$. Firms with deep jurisdictional capital from established EU market presence are better positioned to navigate and implement the Act's requirements, and this effect is more salient for high AI firms, for whom the Act is more relevant.

Cumulative abnormal return $\text{CAR}_i[\tau_1, \tau_2] = \sum_{\tau=\tau_1}^{\tau_2} \widehat{\epsilon}_{i\tau}$ is the sum of daily abnormal returns $\widehat{\epsilon}_{i\tau}$ over the event window. Following Fama–French three-factor model (Fama and French, 1993), daily abnormal return $\widehat{\epsilon}_{i\tau}$ is the residual from the following regression

$$r_{it} - r_{ft} = \alpha_i + \eta_{1i}(\text{Mkt-RF})_t + \eta_{2i}\text{SMB}_t + \eta_{3i}\text{HML}_t + \varepsilon_{it},$$

estimated over trading days $[-250, -11]$ relative to the event with a minimum of 120 valid days (median in-sample $R^2 = 0.34$). We report four windows: $[0, +2]$ as primary, $[0, +1]$, $[-1, +1]$, and $[-2, +2]$ as robustness.

For the multi-event analysis, we re-estimate CARs around each subsequent legislative milestone using the same three-factor methodology and hold the exposure measures fixed at their April 2021 values. This predetermined-exposure design ensures that cross-event variation in coefficients reflects differences in the information content of each event, not changes in firm characteristics over the 2021–2024 legislative period.

5 Main Results

5.1 Baseline Results

Table 1 (Panel A) reports the estimates (Equation 1) across the four event windows. The EU×AI Growth interaction is positive and significant at the 1% level in every window: +40 bps on CAR$[0, +2]$, with comparable magnitudes on CAR$[0, +1]$, CAR$[-1, +1]$, and CAR$[-2, +2]$. The +39.7-basis-point interaction coefficient on $CAR[0, +2]$ implies that a one-standard-deviation increase in pre-event AI hiring growth raises the marginal effect of EU revenue share by 39.7 basis points. Because AI hiring growth captures firms' exposure to the Act through their future AI deployment options, the result indicates that deeper EU market presence was valued more positively among firms more exposed to the new regulatory regime.

The standalone AI hiring growth effect is significantly negative, however, the subsequent checks (see Section 5.5 later) show that the effect is not robust to, for instance, longer term effects (e.g., Table A8) or macro news sensitivity Table A14. The standalone EU revenue share effect is economically small and not consistently significant across windows, confirming that the result largely operates through the interaction of EU commitment and AI activity.

Table 1 (Panel B) traces how the marginal EU effect varies with AI growth, evaluating $\partial \text{CAR}/\partial \text{EU} = \beta_3 + \beta_1 AI$ at three points in the AI-growth distribution. At mean AI growth, the estimated marginal effect of EU exposure is close to zero; at one standard deviation above mean AI growth, it rises to +42.1 basis points. This pattern indicates that EU market presence carried no general announcement premium; its valuation effect became more positive as firms' exposure to the Act increased.

Economic magnitude. To scale the announcement-day repricing to aggregate market capitalization, we compute the dollar change ($\Delta\text{MktCap} = \text{MktCap}_{\text{Apr20}} \times \text{CAR}[0, +2]$) for each firm and aggregate to the total impact. The high EU × high AI firms (i.e., above median in both dimensions) lost only 9 bps (−$16 billion loss on a $17.5 trillion base), while low EU × high AI firms lost 58 bps (−$47 billion on a $8.1 trillion base). The 49 bps differential illustrates the market valuation of jurisdictional capital. Applying the 49 bps differential to the high EU group's base implies jurisdictional commitment shielded roughly $86 billion (49bps×$17.5 trillion base) in shareholder value at announcement.

Other legislative events. We trace the EU×AI interaction across the four legislative milestones of the AI Act, holding the sample fixed at the April 2021 firms so that cross-event variation in coefficients reflects information content rather than composition. The interaction

Table 1: Baseline: EU×AI and Abnormal Returns

Panel A. Baseline regressions

	Dependent variable: CAR (bps)			
	(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
EU × AI Growth	+31.8***	+39.7***	+43.4***	+45.2***
	(9.0)	(10.3)	(7.2)	(8.8)
EU Revenue Share	+3.2	+2.4	+15.1*	+22.5**
	(9.4)	(13.9)	(8.5)	(11.1)
AI Growth	-21.9***	-16.7**	-37.1***	-32.2***
	(6.0)	(6.4)	(8.0)	(7.6)
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Observations	769	769	769	769

Panel B. Marginal effect of EU exposure on CAR$[0,+2]$

AI growth	Marginal EU effect	FF48-clustered p	Wild-cluster p	N
1 SD below mean	-37.3 bps	0.092	0.163	769
At mean	+2.4 bps	0.862	0.897	769
1 SD above mean	+42.1 bps	0.0005	0.053	769

Notes. Panel A reports OLS regressions of cumulative abnormal returns (CAR, in bps) around the 21 April 2021 EU AI Act proposal, based on Equation 1. Standard errors are clustered at the FF48 industry level and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Panel B evaluates $\partial CAR/\partial EU = \beta_{EU} + \beta_{EU\times AI} AI$ using the CAR$[0,+2]$ regression and its full clustered covariance matrix.

is large and significant at the 1% level only at the April 2021 Commission proposal—the moment at which the Act's risk-based architecture, extraterritorial scope, and enforcement regime were first revealed—and indistinguishable from zero at every subsequent milestone; the AI growth standalone effect follows the same pattern. Full results are in Appendix A5.

Another related event is the leakage of the Act draft on April 14, 2021,[7] and civil-society groups had also responded publicly before the official release discussing the leaked draft. As shown in Table A6, both the leak-only short window and leak-to-announcement window—CAR$[-5,-3]$ and CAR$[-5,-1]$, respectively—have a positive but insignificant EU×AI coefficient. Combining the leak and official-release periods, CAR$[-5,+2]$ gives a significant 54.3 bps EU×AI coefficient. These results clarify the role of treatment timing. If some investors priced the leaked draft before April 21, the primary CAR$[0,+2]$ estimate should be interpreted as the official-release component of the market response rather than the total information effect. The full CAR$[-5,+2]$ window captures the broader short-window response from the first public leak date through two trading days after the official publication. The combined-window estimate is larger than the official-window estimate and preserves the same cross-sectional pattern, so leakage makes the official-window magnitude

[7] https://www.politico.eu/wp-content/uploads/2021/04/14/AI-Draft.pdf

conservative rather than overturning the interpretation.

5.2 Pre-Trend Test

The identifying assumption is that, absent the EU AI Act announcement, high EU×AI firms would have experienced abnormal returns similar to other firms. We test this using a symmetric 81-day window $[-40, +40]$ around the event. We shift the abnormal return estimation window back by 40 days relative to the main specification: factor loadings are estimated over $[-291, -41]$ (250 trading days, same minimum 120-day coverage requirement), so the pre-trend window $[-40, -1]$ is fully disjoint from the CAR estimation window and the test is not mechanically contaminated by in-sample fit.

Table 2 reports the EU×AI coefficient across pre-event cumulative windows. The coefficient is economically small and statistically indistinguishable from zero in every pre-event window, in contrast to the significant event-window coefficients post-event. The effect emerges precisely at the event date, not before. Appendix Figure A.1 shows the full cumulative coefficient trajectory from day −40 through day +40: the path is mostly flat and near zero throughout the pre-event period, then rises sharply after day 0 and stabilizes at approximately +100 bps by day +40.

Table 2: Pre-Trend Test: EU×AI Coefficient Across Cumulative Windows

Window	EU × AI (bps)	SE (bps)	t	p-value	N
Pre-event					
CAR[-40,-1]	-25.2	59.5	-0.42	0.674	769
CAR[-40,-21]	-10.9	41.8	-0.26	0.795	769
CAR[-20,-1]	-14.3	23.4	-0.61	0.545	769
CAR[-10,-1]	4.3	18.6	0.23	0.818	769
CAR[-5,-1]	14.6	11.3	1.29	0.205	769
Event / post-event					
CAR[0,+2]	39.7***	10.3	3.87	<0.001	769
CAR[0,+10]	86.7***	32.2	2.70	0.010	769
CAR[0,+40]	174.0***	54.4	3.20	0.003	769

Notes. Each row reports the EU×AI coefficient from a separate OLS regression of the indicated cumulative abnormal return (CAR) window on EU×AI, EU Revenue Share, AI Growth, firm controls, and FF48 industry fixed effects, using the baseline specification (N=769 after singleton FF48-cluster drop). Standard errors clustered at the FF48 industry level. *, **, *** denote significance at the 10%, 5%, 1% levels. Pre-event windows are placebos; their imprecise coefficients provide no evidence of differential pre-trends.

5.3 Randomization Inference

Randomization inference provides a distribution-free check on our standard-error computation by comparing the observed coefficient to a placebo distribution built from non-event trading dates. We treat all 483 trading dates between December 2020 and December 2022 as placebo event dates, excluding a ±10-day buffer around April 21, 2021 to avoid mechanical overlap between any placebo window and the actual event window. For each placebo date, we recompute CARs using an estimation window of $[-250, -11]$ trading days relative to the placebo and re-estimate the full interaction model with identical specification.

The observed April 2021 coefficient on the primary CAR$[0, +2]$ window lies in the far right tail of the placebo distribution for the EU×AI interaction (Figure 4). Both coefficients are rejected by placebo inference at conventional significance levels. In Appendix Table A7 and Appendix Figures A.2–A.3, we show coefficients for the full set of windows where they are significant on all four windows. The EU revenue share only effect, in contrast, falls well within its placebo distribution, reinforcing that the result operates through the interaction rather than through EU exposure alone.

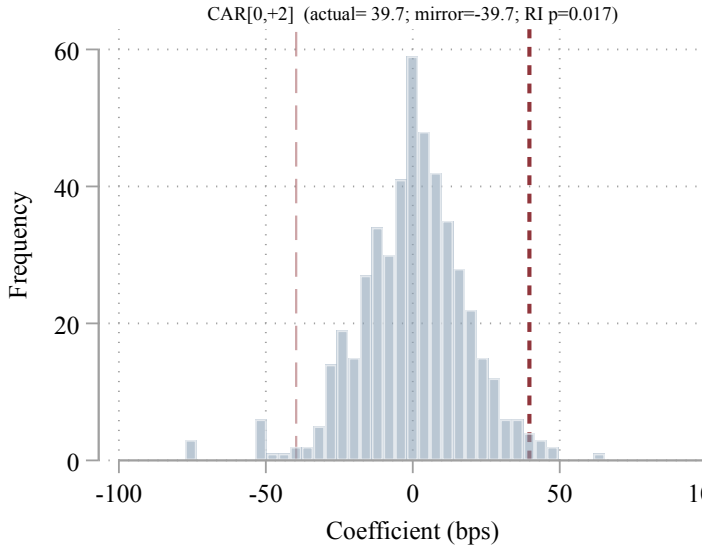


Figure 4: Randomization Inference on CAR$[0, +2]$: Actual Coefficients vs. Placebo Distributions

Notes: The figure plots the distribution of coefficients of EU×AI from 483 placebo regressions run on non-event trading dates between December 2020 and December 2022, excluding a ±10-day buffer around April 21, 2021. The solid red line marks the actual April 21, 2021 coefficient. Full 2×2 distributions across all four event windows are in Appendix Figures A.2–A.3.

5.4 AI Growth vs. AI Level

A natural question about our specification is if the jurisdictional capital effect operates through pre-existing AI capability or through the change in capability before the Act's proposal. The level of AI hiring intensity, after all, is a more direct proxy for current AI exposure than the growth rate, and a stock-based reading of the regulatory channel might predict that firms with more AI activity at the time of the proposal needs more adaption to the Act. We address this concern through two tests: (i) replacing AI hiring growth with the AI hiring level in the baseline specification, and (ii) running a horse race that enters both measures simultaneously.

Table 3: Horse Race: EU×AI Level vs. EU×AI Growth

	Dependent variable: CAR (bps)			
	(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
EU × AI Level	-9.1	0.2	-14.3	1.2
	(7.7)	(14.4)	(16.2)	(12.8)
EU × AI Growth	37.2**	39.6**	52.0***	44.5*
	(14.7)	(17.8)	(18.1)	(23.7)
AI Level	11.4	4.9	8.6	3.1
	(13.2)	(15.9)	(16.3)	(21.3)
AI Growth	-25.3**	-18.2	-39.6***	-33.2*
	(11.4)	(13.8)	(14.1)	(18.4)
Joint F: EU × (Level+Growth)=0	7.20	9.58	20.60	14.84
	$p = 0.002$	$p \approx 0.000$	$p \approx 0.000$	$p \approx 0.000$
Industry FE	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
N	769	769	769	769
Adj. R^2	0.017	0.086	0.050	0.077

Notes. This table extends Table 1 by including AI Level and its interaction with EU share. Standard errors clustered at the FF48 industry level are in parentheses. The joint F-test reports the null $\beta(\text{EU} \times \text{AI Level}) + \beta(\text{EU} \times \text{AI Growth}) = 0$. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

AI level doesn't exhibit mitigating effects when interacted with EU share. As shown in Table 3, the coefficient of AI level is not significant in any of the event windows, and the coefficient of EU×AI Level is approximately half the magnitude of the EU×AI growth in Table 1 and not robust.

A more demanding test enters EU×AI level and EU×AI growth simultaneously, alongside AI main effects (Table 3). The EU×AI growth coefficient is essentially unchanged from the headline (+40 bps on CAR$[0,+2]$) and remains significant on all four event windows. The EU×AI level coefficient is insignificant in every window, with a point estimate of +0 bps

on CAR$[0, +2]$. This is not a multicollinearity artifact: the two interactions correlate at $\rho = 0.60$ in the regression sample, well below thresholds at which a real signal would be masked.

Why growth, not level. The horse race confirms that the Act operates through the growth margin: pre-existing AI intensity carries no incremental EU-interactive return reaction once the growth in AI capability is in the model. This pattern is consistent with the Act regulating AI systems placed on the EU market going forward: its compliance obligations bind on future AI footprint rather than on the firm's existing AI footprint. A firm with high static AI intensity but no further expansion has, in expectation, already amortized its AI investment and faces lower marginal compliance cost from new deployment. A firm scaling its AI capabilities rapidly during the pre-event window has growth options that are repriced when the regulatory regime is announced, because those growth options will be deployed under the new compliance regime. The effects therefore load on growth, not level—the market prices the regulation according to firm's future AI capabilities based on current investment trajectory, rather than merely the AI capability already in place.

5.5 Other Robustness Checks

We test if the effect persists over longer horizon by extending the window of CAR. The results suggest that the jurisdictional capital effect from EU market presence is maintained through several weeks, providing evidence against short-lived overreaction (Appendix B.3 and B.4).

Other robustness checks are replacing the Fama–French three-factor model with the five-factor specification (Appendix B.5), replacing EU revenue share with other pre-2021 years or including 2020, using the firms whose fiscal year ends in December (Appendix B.6),[8] add controls for firm's equity duration and sensitivity to general macro news (Appendix B.7), excluding potential tax-haven countries in EU revenue share,[9] excluding April 22 when the US administration released a capital gains tax proposal,[10] and removing potential outliers.[11]

[8]Two features of our data construction dilute potential fiscal year and calendar year misalignment ex-ante: we average each input over five years (FY/CY 2015–2019) and we standardize within sample after the average.

[9]Excluding Ireland, Luxembourg, and the Netherlands from EU share yields 41.0 bp (p-value = 0.082).

[10]Within our CAR[0,+2] window, on April 22, the US administration proposed nearly doubling the top federal capital-gains tax rate on high-income individual investors, triggering a broad stock market selloff. The results nonetheless remains broadly unchanged after controlling for equity duration and sensitivity to yield and VIX news, which are the characteristics most likely to transmit such a shock. Excluding April 22, the remaining days remain both positive and significant.

[11]Both AI hiring growth and EU revenue share are right-skewed, so we re-estimate the interaction under several alternatives. Deleting the five and ten largest DFBETA observations yields 37.2 and 34.5 bp (p-value

None of them overturn statistical inference. Lastly, we provide evidence consistent with the jurisdictional reach of the AI, showing that the proposal announcement only affects the firms with EU revenue (Appendix B.8).

6 Competing Interpretations

The baseline EU×AI interaction establishes that firms with deeper EU market commitment while growing AI capabilities earned positive abnormal returns around the AI Act proposal. Our story is that jurisdictional capital—accumulated compliance infrastructure and institutional knowledge from years of operating under EU regulation—helps firms navigate the AI Act's requirements. However, other competing interpretations could also be consistent with this finding. For instance, the EU×AI interaction may represent other firm characteristics—firm size, international diversification—rather than jurisdictional capital. Or it could reflect lobby effects: firms deeply rooted in the EU may have lobbied to shape the Act in their favor.

This section tests these alternative explanations. Section 6.1 distinguishes jurisdictional commitment from firm scale, separating our finding from the scale-based advantage documented under GDPR. Section 6.2 then addresses regulatory capture by testing political-access channels directly using various EU and member state level lobby proxies. Section 6.3 confirms that the effect is specific to EU jurisdiction rather than general international diversification.

6.1 Jurisdictional Commitment, Not Scale

The existing literature on technology regulation emphasizes scale-based benefit: large firms are less negatively affected by regulation than small players, because, for instance in the case of GDPR, compliance fixed costs could be amortized over a larger revenue base. Our EU revenue share measure could capture just the scale of the firm's EU operations rather than anything about jurisdictional capital. To test whether the jurisdictional capital effect operates through the same scale-based mechanism, we modify the baseline specification by replacing EU revenue share with alternative scale proxies as the interaction variable and testing whether each generates a comparable effect.

Table 4, Panel A replaces EU revenue share with absolute EU revenue in USD and its logarithm, which are direct measures of the absolute size of a firm's EU operations.

$= 0.038, 0.016$), so the result is not outlier-driven. Estimates remain positive but attenuate under 1/99 winsorizing (33.4, p-value = 0.066) and posting floors of 25 postings (32.2, p-value = 0.078).

Table 4: Jurisdictional Commitment vs. Scale

Variable	Dependent variable: CAR (bps)			
	(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
Panel A: Alternative scale measures of EU exposure × AI Growth				
EU Share × AI Growth	31.8***	39.7***	43.4***	45.2***
	(9.0)	(10.3)	(7.2)	(8.8)
EU USD × AI Growth	-0.9	-2.8	8.3	5.4
	(4.0)	(4.6)	(7.0)	(7.0)
log(EU USD) × AI Growth	1.4	-4.4	14.2	4.6
	(8.2)	(10.0)	(10.3)	(11.7)
Panel B: Size × AI Growth horse race against the baseline EU × AI				
EU × AI Growth	33.5***	41.5***	45.4***	47.6***
	(9.4)	(10.7)	(7.4)	(9.2)
log(Mcap) × AI Growth	-12.8	-13.4	-14.5*	-17.6*
	(7.5)	(9.1)	(8.3)	(11.4)
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls (incl. lower-order interactions)	Yes	Yes	Yes	Yes
Observations	769	769	769	769

Notes. Panel A replaces EU Revenue Share with three alternative scale measures of EU exposure: (i) the baseline EU revenue share; (ii) the corresponding raw EU revenue in USD; and (iii) log(1 + EU revenue USD). Panel B adds log(Mcap) × AI Growth to the baseline EU × AI specification (all lower-order terms included). All specifications include FF48 industry fixed effects and the five firm controls of Equation (1). FF48-clustered SE in parentheses. *, **, *** denote significance at 10%, 5%, 1%.

The interaction between absolute EU revenue and AI growth is economically small and statistically insignificant across all event windows. The jurisdictional capital effect does not emerge from the absolute size of EU operations; it requires the relative importance of the EU market to the firm's total business. Panel B shows a horse race with both EU share and firm size interacting with AI growth, the coefficient for firm size×AI is small in magnitude and low in significance. Meanwhile, EU×AI remains significant and close to baseline estimates.

Table 5 extends the test in Panel B to three additional firm-level scale proxies, each capturing a distinct dimension along which scale could plausibly operate, estimated on CAR$[0,+2]$. General size—log employment (column 4) alongside log market capitalization (column 2)—carries a negative interaction with AI growth (around −13 bps), suggesting larger firms bearing a modest compliance penalty on their AI operations. Physical capital intensity, measured as PP&E to total assets (column 3), loads strongly negative (−49 bps): firms whose AI systems are embedded in physical capital—e.g., manufacturing, telecom, utilities—face the heaviest compliance requirements. It could be that AI embedded in physical production processes is operationally more costly to modify than software-only deployments. R&D intensity, a proxy for the scale of AI development investment, is essentially zero. In every specification, the EU×AI effect remains significant at the 1% level.

Row 6 of Table 5 enters EU×AI jointly with all three scale interactions. The general-size loading attenuates once physical capital intensity enters, while the physical-capital-intensity penalty remains highly significant; the EU×AI effect holds and the R&D interaction remains null. The EU jurisdictional capital effect therefore operates through a dimension different from four distinct dimensions of firm scale—absolute EU operations, general size, physical capital intensity, and R&D intensity.

Table 5: Scale Decomposition: EU Revenue Share vs. Alternative Firm-Level Proxies

	Coefficient on ·× AI Growth (bps), CAR[0, +2]					
Variable	(1) Baseline	(2) +log(Mcap)	(3) +PP&E/AT	(4) +log(Employ)	(5) +R&D	(6) Saturated
EU × AI	39.7***	41.5***	34.0***	42.1***	39.7***	35.4***
	(10.3)	(10.7)	(9.5)	(11.3)	(10.3)	(9.4)
log(Mcap)	—	-13.4	—	—	—	-10.4
		(9.1)				(8.7)
PP&E/AT	—	—	-48.8***	—	—	-47.5***
			(14.6)			(15.5)
log(Employ)	—	—	—	-12.8*	—	—
				(7.1)		
R&D Intensity	—	—	—	—	1.5	-4.8
					(8.1)	(5.2)
Industry FE (FF48)	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	769	769	769	769	769	769

Notes. Each column reports a regression of CAR[0, +2] on EU × AI Growth plus the cost-structure interaction indicated in the column header, with all five baseline firm controls and FF48 industry fixed effects. log(Mcap) is the natural log of market capitalization in the month prior to the event date. PP&E/AT is net property, plant, and equipment divided by book total assets. log(Employ) is the natural log of total employees, proxy for workforce scale. R&D Intensity is R&D expenditure divided by total assets, proxy for the variable cost of developing AI capabilities. FF48-clustered standard errors in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels on the cluster-robust t-statistic.

GDPR vs. AI Act. The contrast between GDPR and the AI Act helps explain why scale alone is unlikely to account for the jurisdictional capital effect. GDPR compliance has a significant firm level fixed cost component. A firm can build a common privacy and consent system, and once users give consent, that consent can often support data use across multiple services offered by the same firm, subject to the stated purposes. Larger user bases can therefore help spread these costs. By contrast, many AI Act requirements apply at the system or use case level, generating more variable compliance costs: a firm with more AI products may face more regulated systems, scaling up the overall compliance burden.

Second, the AI Act may create a more complex compliance navigation problem than GDPR. GDPR was a major regulatory change, but it revised an existing EU privacy regime,

the 1995 Data Protection Directive, giving firms familiar legal and institutional reference points.[12] Regulators, guidance, and the compliance advisory ecosystem was also relatively mature by the time GDPR was introduced. The AI Act is different. As the world's first comprehensive AI regulatory framework, it creates an uncertain and evolving regime, with many practical details to be clarified gradually by regulators, certification bodies, and standard setters. Stable EU presence may therefore matter because it reflects more than market size. Firms with sustained EU operations are more likely to have local legal capacity, experience with EU regulatory processes, and familiarity with how compliance is implemented in practice. Such prior commitment can lower the cost of adapting to a new regime, even when firm size alone does not.

6.2 Political Engagement and Lobby Effect

Another concern is that firms with deeper EU roots shaped the AI Act in their favor, so the positive EU×AI coefficient reflects the payoff from political influence rather than the value of jurisdictional capital. Firms with deeper EU roots may be better positioned to influence the Act's design—for example, by engaging with legislative bodies, shaping technical standards, or securing favorable interpretations of vague compliance requirements. Before turning to directly test the lobby effect, it is worth noting that as shown in Table 4 Panel B, the Size×AI Growth interaction is small and mostly insignificant: the effect does not vary with firm size, even though political access is strongly correlated with size.

We construct various indicators capturing both the intensive and extensive margins of political engagement at EU and member state levels: (i) *Commissioner meetings*, an indicator for any meeting during 2015–2019 (73 of 771 firms); (ii) *Lobby registration*, an indicator for pre-event appearance in the EU Transparency Register (74 firms); (iii) *AI White Paper consultation*, an indicator for submitted comments in AI White Paper consultation stage (17 firms); (iv) *Country-level lobby*, an indicator equal to one if the firm appears pre-event in any mandatory member-state political-engagement channels: France, Ireland, and Germany (86 firms); (v) *EU-level trade-association membership*, an indicator if the firm appears on the pre-event member roster of DigitalEurope, DOT Europe, or ITI—the three major EU-level digital-industry trade associations (45 firms); and (vi) *Any political engagement*, equal to one if one of the condition holds. Details of the variable construction is in Appendix A.

We run three progressively demanding tests with each political engagement proxy. *Control* adds political engagement as a control to the baseline specification. Under the lobby effect, as political engagement is positively correlated with EU share and abnormal returns,

[12] https://eur-lex.europa.eu/legal-content/en/ALL/?uri=CELEX:31995L0046

the coefficient on EU × AI should attenuate once political engagement is absorbed. *Horse race* includes Lobby×AI alongside EU×AI, and *triple interaction* tests whether political access amplifies the effect through Lobby×EU×AI. Under lobby effect, Lobby×AI coefficient and the triple coefficient should be positive, as political engagement should confer a similar effect in adapting to the Act.

Minimal impact from political engagement variables. Table 6 shows that the EU×AI coefficient is virtually unchanged from its baseline value of +40 bps on CAR$[0, +2]$ when each political engagement indicator is added as a control, as horse race, or as triple interaction. None of the political engagement variables—dummy, interaction term with AI, or triple interaction—show consistent significance. Together, these suggest that political engagement does not substitute for EU presence in generating the jurisdictional capital effect. Besides the results based on $CAR[0, +2]$, the other outcome variables also remain qualitatively unchanged in the three sets of tests.

To complement the analysis on pre-event lobbying, we also examine *post*-April-2021 AI-Act lobbying as an outcome. Across EU-level firm lobbying, country-level lobbying, EU trade-association membership, and their union, the EU×AI interaction does not predict additional lobbying (all effects statistically insignificant), while AI hiring growth and firm size strongly do. Appendix C details the construction of the tests and the full regression evidence.

6.3 Foreign Revenue Share vs. EU Revenue Share

Another potential competing explanation is that the EU×AI interaction captures the benefits of general international diversification rather than EU-specific jurisdictional capital. If so, the interaction of other foreign revenue share with AI should yield a similar positive coefficient.

Table 7 reports the test. We add Foreign×AI Growth to the baseline specification. The EU×AI interaction retains its magnitude and significance across all windows, while the Foreign×AI interaction is economically small and statistically insignificant. General foreign exposure does not generate the jurisdictional capital effect, and only EU-specific market presence does. This rules out the diversification alternative and confirms that the baseline EU×AI interaction captures a channel specific to the EU AI Act's jurisdictional reach.

Table 6: Lobby Effect vs. Jurisdictional Capital Effect

	Dependent variable: CAR$[0,+2]$					
	Engagement = Commissioner meetings			*Engagement = Transparency Register*		
	Control	Horse race	Triple interaction	Control	Horse race	Triple interaction
EU × AI Growth	38.6***	40.0***	43.2***	39.4***	39.5***	39.5***
	(9.8)	(9.9)	(10.7)	(9.7)	(10.1)	(10.3)
Engagement	40.2			43.5*		
	(25.7)			(23.3)		
Engagement × AI		-9.1			-5.1	
		(6.5)			(4.8)	
Engagement × EU × AI			10.2			0.3
			(10.1)			(9.5)
	Engagement = AI White Paper consultation			*Engagement = Country-level lobby*		
	Control	Horse race	Triple interaction	Control	Horse race	Triple interaction
EU × AI Growth	39.6***	39.6***	42.1***	42.2***	41.7***	41.4***
	(10.2)	(10.0)	(9.7)	(10.4)	(9.9)	(8.9)
Engagement	61.3			-122.6**		
	(105.2)			(49.0)		
Engagement × AI		3.9			2.6	
		(4.6)			(7.8)	
Engagement × EU × AI			-7.1			2.5
			(9.8)			(11.3)
	Engagement = EU-level trade association			*Any political engagement*		
	Control	Horse race	Triple interaction	Control	Horse race	Triple interaction
EU × AI Growth	39.7***	39.9***	41.1***	39.5***	39.5***	38.5***
	(10.4)	(10.9)	(10.2)	(10.2)	(10.1)	(10.7)
Engagement	-3.9			13.6		
	(49.4)			(33.8)		
Engagement × AI		-20.1			-2.4	
		(13.3)			(13.3)	
Engagement × EU × AI			-17.8			-12.1
			(13.2)			(17.7)
Industry FE (FF48)	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	769	769	769	769	769	769

Notes. This table reports OLS regressions of a regression of CAR$[0,+2]$ around the 21 April 2021 EU AI Act proposal, augmenting baseline specification Equation 1 with political engagement dummy variables as additional control, horse race, and triple interaction. Standard errors are clustered at the FF48 industry level and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table 7: EU vs. Foreign Revenue Share

	Dependent variable: CAR $[0,+2]$ (bps)			
	(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
EU × AI Growth	29.6***	34.0***	44.6***	43.0***
	(10.2)	(12.0)	(7.8)	(7.9)
Foreign Income × AI	1.8	0.3	4.8	2.1
	(7.8)	(10.6)	(8.7)	(13.7)
EU Revenue Share	5.2	16.4*	19.0**	42.8***
	(6.5)	(9.5)	(9.4)	(10.1)
Foreign Income	-12.6	-10.7	-20.3*	-17.7
	(9.2)	(13.5)	(10.8)	(13.0)
AI Growth	-22.4***	-22.4***	-33.4***	-32.8***
	(6.4)	(7.9)	(8.9)	(8.6)
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Observations	681	681	681	681
Adjusted R^2	0.023	0.042	0.057	0.021

Notes. This table reports OLS regressions of a regression of CAR$[0,+2]$ in bps around the 21 April 2021 EU AI Act proposal, with foreign-income interaction (non-US, non-EU revenue share). Standard errors are clustered at the FF48 industry level and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

7 Amplifiers of the Jurisdictional Capital Effect

The baseline results establish that EU×AI exposure predicts positive returns on average. We now test whether the jurisdictional capital effect is amplified by observable firm characteristics: high-risk AI capability buildup (Section 7.1), stability of EU market presence (Section 7.2), and geographic concentration of EU operations (Section 7.3). Section 7.4 provides a more direct test of the mechanism by asking whether pre-event compliance hiring, an observable proxy for regulatory capital accumulation, amplifies the effect independently of EU revenue share.

7.1 The Jurisdictional Capital Effect Is Stronger in High-Risk AI

The Act's most stringent requirements—conformity assessments, Annex IV technical documentation, and post-market monitoring under Articles 16–22—attach specifically to systems listed in Annex III. If the jurisdictional capital helps firms navigate the Act's requirement, it should be largest where the regulatory requirements are greatest: firms scaling up into high-risk-AI involvement. We test this using a dictionary-based mapping of Lightcast skills and titles to the Annex III scope.

We construct a high-risk-AI hiring measure by using the skill and title information from

Lightcast postings. A posting is classified as high-risk AI if it jointly matches our high-risk-AI skill dictionary and our high-risk-AI title dictionary, constructed by mapping the Lightcast universe to Annex III categories (biometric identification, education and vocational training, employment and worker management, creditworthiness assessment) and Article 6(1) Annex II Section A safety components (robotics and medical-device AI). Deployer-restricted Annex III categories, safety-component carve-outs, and generic HR/recruiting patterns that flag at essentially every firm are excluded. HRAI_i is the 2015–2019 annualised change in firm i's share of postings so classified. A firm scaling into high-risk-AI hiring during the pre-event window is actively building capabilities in the domains Annex III and Annex II Section A target, making it more likely to produce or deploy high-risk AI systems at the time of the proposal and therefore to face the Act's most demanding compliance obligations. Full construction is in Appendix D.

Panel A of Table 8 shows that the jurisdictional capital effect grows with firm's high-risk-AI hiring: on the CAR$[0, +2]$ window, each additional standard deviation of HRAI_i adds approximately +15 bps to the baseline effect, a pattern that holds across all four event windows, while the baseline EU$\times$AI coefficient remains positive. Panel B makes the economic content concrete: splitting the sample by whether firms have any high-risk-AI hiring growth, the effect is large and statistically significant in the high-risk-AI group and economically negligible in its complement, with the contrast widening sharply at longer horizons. EU market presence alone is insufficient to generate the full jurisdictional capital effect.

The results identify where jurisdictional capital is most valuable and why. The effect amplifies with high-risk-AI hiring because these systems carry the Act's heaviest and most interpretively complex obligations, so jurisdictional capital is worth most precisely where regulatory demands are greatest. The binary subsamples sharpen this reading: the high-risk subsample shows a large positive premium (+35 bps on CAR$[0, +2]$), while the zero-HRAI subsample yields an insignificant estimate, indicating that EU presence alone does not generate the effect. Together, they imply that the baseline EU$\times$AI coefficient in Section 5.1 is driven by firms at the intersection of EU commitment and high-risk-AI exposure, where both jurisdictional and regulatory demands bind most tightly.

7.2 Stable EU Commitment Matters

This section tests if the stability of EU revenue commitment, measured by the coefficient of variation (CV, i.e., standard deviation divided by mean) of EU revenue share over time, matters for the magnitude of the jurisdictional capital effect. Intuitively, stable EU presence—low year-to-year fluctuation in EU share—signals sustained strategic commitment to the EU

Table 8: High-Risk AI

	Dependent variable: CAR (bps)			
	(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
Panel A: Triple interaction, full sample				
EU × AI × HRAI	10.18***	15.02***	15.82***	20.58***
	(2.95)	(2.81)	(3.08)	(3.33)
EU × AI	25.53*	31.69**	26.57**	23.56
	(1.83)	(2.03)	(2.08)	(1.63)
Panel B: Binary subsample split by HRAI hiring growth				
EU × AI, HRAI > 0 ($N = 202$)	30.76***	34.86***	72.17***	83.65***
	(4.07)	(3.46)	(6.95)	(4.47)
EU × AI, HRAI $= 0$ ($N = 554$)	27.94	27.83	12.55	1.52
	(1.35)	(1.49)	(0.42)	(0.05)
Firm controls	Yes	Yes	Yes	Yes
Industry FE (FF48)	Yes	Yes	Yes	Yes
Cluster	FF48	FF48	FF48	FF48

Notes. Panel A reports the triple-interaction coefficient $\hat{\beta}_1$ on $\text{EU}_i \times \text{AI}_i \times \text{HRAI}_i$ and the EU×AI two-way coefficient $\hat{\beta}_2$ from the same regression; all three continuous variables are in-sample z-scores. Panel B re-estimates the baseline EU×AI specification within subsamples defined by whether the firm has any non-zero high-risk-AI hiring growth over 2015–2019; treatment variables are re-standardised within each subsample. Standard errors clustered on FF48 industry. *, **, *** denote 10%, 5%, 1%, respectively.

market, while volatile EU share suggests opportunistic or episodic engagement. This strategic commitment, rather than episodic engagement, would accumulate jurisdictional capital over time.

We split the sample at the median of CV and re-estimate the baseline regression. The prediction is that the coefficient based on the low-CV sample should be larger: firms with more stable EU presence (lower CV) should have a stronger effect, because stability reflects accumulated jurisdictional capital rather than transient market participation. Table 9 reports the results. The coefficient of EU×AI for the low-CV sample is significantly positive and close to full sample estimate, however, its high-CV sample counterpart is insignificant across all four event windows.

7.3 Geographic Concentration of EU Revenue

The jurisdictional capital effect hinges on regulatory knowledge and compliance infrastructure buildup. As the implementation falls at member state level, the jurisdictional capital also crucially depends on tacit, country-specific enforcement knowledge—e.g., how rules are applied in practice and what regulators expect. Holding total EU revenue share fixed, this

Table 9: EU Revenue Stability

		Dependent variable: CAR (bps)			
		(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
Panel A: Low-CV subsample (stable EU presence, below-median CV)					
Low-CV	EU × AI Growth	35.0***	38.0***	58.9***	56.7***
		(10.1)	(8.7)	(8.9)	(9.4)
N				371	
Panel B: High-CV subsample (volatile EU presence, above-median CV)					
High-CV	EU × AI Growth	-4.5	12.9	-13.3	-4.0
		(14.6)	(22.4)	(22.7)	(35.9)
N				368	
Industry FE (FF48)		Yes	Yes	Yes	Yes
Firm controls		Yes	Yes	Yes	Yes

Notes. This table reports OLS regressions of cumulative abnormal returns (CAR, in bps) around the 21 April 2021 EU AI Act proposal, estimated separately on below-median and above-median subsamples of the firm-level coefficient of variation (CV) of the EU revenue share over FY2010–2019. CV is the standard deviation of the firm's annual EU share divided by its mean over the window, winsorized 1%/99% within the analysis sample. Standard errors, clustered at the FF48 industry level, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

capital should be larger when revenue is concentrated in a few member states than when it is dispersed. For example, 10% revenue concentrated in Germany builds depth with German enforcement institutions (e.g., *Datenschutzkonferenz*, *BfDI*), whereas 10% spread thinly across many member states is unlikely to generate comparable depth anywhere; under an EU-wide regime like the AI Act, the former converts this depth into benefit. We test this mechanism using within-EU revenue concentration: the Herfindahl index (HHI)[13] across member states and the top-1 member-state revenue share.

Table 10 splits the sample at the median of each measure and re-estimates the baseline regression. Under the HHI measure (Panel A), the high-concentration half shows a strong EU×AI coefficient across all four event windows, ranging from +47 bps to +86 bps, all significant at the 1% level. The low-concentration half is statistically and economically null on three of the four windows. Panel B (top-1 split) delivers a similar pattern. The magnitude gap between high- and low-concentration subsamples widens from roughly six× on the three-day window (column 2) to 20× on the five-day symmetric window (column 4). This supports

[13]For each firm f, define the within-EU share of member state c as $w_{f,c} = s_{f,c}/S_f$, where $s_{f,c}$ is the revenue from member c and S_f is the total EU revenue. The within-EU Herfindahl index $H_f = \sum_c w_{f,c}^2$, which ranges from $1/K$ (perfect diversification across K disclosed countries) to 1 (a single country carries the entire EU exposure).

the mechanism: depth of engagement, not breadth, converts EU market participation into a regulatory know-how. Two firms with identical total EU revenue share can face entirely different event returns depending on whether that revenue is anchored in specific member states or widely dispersed across the EU.

Table 10: Geographic Concentration of EU Revenue

		Dependent variable: CAR (bps)			
		(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
Panel A: Split by HHI (within-EU Herfindahl)					
High-concentration (above median)	EU × AI Growth	47.0***	63.9***	75.1***	86.0***
		(13.9)	(19.7)	(12.3)	(15.4)
N				382	
Low-concentration (below median)	EU × AI Growth	21.0**	11.4	13.8	5.9
		(7.8)	(11.5)	(13.1)	(17.1)
N				375	
Panel B: Split by Top-1 within-EU country share					
High-concentration (above median)	EU × AI Growth	47.2***	64.5***	75.3***	87.3***
		(13.3)	(18.7)	(12.1)	(14.6)
N				381	
Low-concentration (below median)	EU × AI Growth	20.9***	9.9	13.0	3.4
		(7.6)	(11.3)	(13.5)	(17.6)
N				376	
Industry FE (FF48)		Yes	Yes	Yes	Yes
Firm controls		Yes	Yes	Yes	Yes

Notes. This table reports OLS regressions of cumulative abnormal returns (CAR, in bps) around the 21 April 2021 EU AI Act proposal, estimated separately on above-median and below-median subsamples of EU concentration. *Panel A* splits the sample at the median of the within-EU HHI. *Panel B* splits the sample at the median of the Top-1 within-EU share. Standard errors, clustered at the FF48 industry level, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

A natural follow-up is whether the concentration-based jurisdictional capital reflects broad capital or capital specific to a single national jurisdiction—for example, a Germany-only or Ireland-only effect that happens to load on the HHI measure. Appendix E decomposes each HHI-based subsample by country of EU revenue and shows that major economies in the EU—Germany, France, Italy, and Spain—dominate both halves, and the high-HHI firms simply concentrate more sharply within that same set. The effect is therefore not a single-country artifact.

7.4 Prior Compliance Experience

The mechanism tests so far take EU revenue share as an indirect proxy for jurisdictional capital, we now test the proposition more directly using firms' pre-event scaling of IT, privacy, and data-governance compliance hiring as an observable proxy for jurisdictional capital accumulation. The intuition is that GDPR's enforcement forced EU-exposed firms to build

compliance capacity and to actively interact with EU regulators. The accumulated compliance infrastructure and interpretative capability are expected to carry over to AI Act obligations.

We identify compliance-related postings through a text analysis over the Lightcast skills and titles dictionaries, applying a curated screen of 77 patterns across 11 functional categories (see more in Appendix F.4 and F). For each firm-year, we then compute the share of postings flagged as compliance-related; firm-level compliance hiring growth is the four-year-annualized change in this share over 2015-2019, constructed in parallel with AI hiring growth. We construct two variants of the measure: a *Broader IT/data compliance* measure that includes US-, UK-, and EU-statute hiring patterns, and an *EU IT/data compliance* measure that strips out the US- and UK-statute matches. The Broader measure captures general compliance capacity, and the EU measure isolates the hiring most plausibly driven by GDPR and adjacent EU regulatory exposure.

We split the sample at the median, re-estimate the baseline regression within each subsample. The EU×AI effect concentrates among compliance-scaling firms (Table 11): in the top half, the effect is approximately +50 bps on CAR$[0, +2]$ and significant at the 1% level; in the bottom half, the effect is largely insignificant. This indicates that jurisdictional capital is prominent in firms with active interactions with regulators through prior IT compliance activities, rather than firms investing less EU regulation compliance infrastructure (e.g., passive EU sales through distributors).

A natural concern is that compliance growth simply relabels the EU revenue channel. In the horse race of Table 12, the EU×AI coefficient is essentially unchanged from baseline while Compliance×AI is indistinguishable from zero on every window. Together with the subsample split, this identifies a moderator structure: stable EU presence and pre-event compliance scaling jointly generate the premium, but neither alone suffices. The EU×AI effect appears only among high-compliance-growth firms, suggesting active EU operations and sustained regulation implementation experiences rather than passive sales through distributors. The null on Compliance×AI is therefore informative: because the AI Act is a new, product- and use-case-based regime, IT-oriented GDPR compliance capacity does not carry over mechanically, and it becomes valuable only when paired with EU-embedded jurisdictional capital.

Table 11: Pre-AI Act Compliance Hiring

		Dependent variable: CAR (bps)			
		(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
Panel A: Broader IT/data compliance growth					
Median: Top half	EU × AI Growth	40.9***	49.5***	59.0***	57.5***
	(N =376)	(11.5)	(12.6)	(11.4)	(9.4)
Median: Bottom half	EU × AI Growth	23.6*	15.4	24.6	15.3
	(N =379)	(12.4)	(16.1)	(26.7)	(26.1)
Panel B: EU IT/data compliance growth (excl. US/UK)					
Median: Top half	EU × AI Growth	44.6***	50.4***	63.4***	60.0***
	(N =378)	(13.6)	(13.9)	(12.5)	(9.7)
Median: Bottom half	EU × AI Growth	19.1	14.4	21.2	10.6
	(N =379)	(12.3)	(15.0)	(24.5)	(24.4)
Industry FE (FF48)		Yes	Yes	Yes	Yes
Firm controls		Yes	Yes	Yes	Yes

Notes. This table reports OLS regressions of cumulative abnormal returns (CAR, in bps) around the 21 April 2021 EU AI Act proposal. *Panel A* cuts the analysis sample by the four-year-annualised growth in the share of postings with at least one core-tier match in our IT/data compliance taxonomy ("Broader" = any core-tier compliance match, including US- and UK-statute categories). *Panel B* cuts by the same growth measure restricted to postings whose core matches are not exclusively in US- or UK-statute categories. Standard errors clustered at FF48 in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

8 EU Jurisdictional Capital and Post-Event Within-Firm Market-Share Reallocation

The results so far demonstrate that jurisdictional capital helps firms navigate and implement the Act's requirements: firms with deep EU-embeddedness adapt more readily than peers lacking such exposure. A natural follow-up question is whether this effect translates into a measurable firm-level benefit beyond the returns reflected in the financial market. One observable margin is the within-firm EU market revenue relative to non-EU market revenue, or similarly, EU revenue share out of global revenue. This margin captures within-firm reallocation of revenue between EU and non-EU regions instead of market share or cross-border entry. We document a cross-sectional pattern consistent with this mechanism: among high EU and high AI firms, those whose peers are lower in EU share, i.e., those find it easier to navigate the Act, experienced systematically increase in the EU share of their revenue after the Act proposal announcement.

Table 12: Horse Race: EU×AI vs. Compliance×AI

		Dependent variable: CAR (bps)			
		(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
Panel A: C = Broader IT/data compliance growth					
(1) Baseline	EU × AI Growth	31.8***	39.7***	43.4***	45.2***
		(9.0)	(10.3)	(7.2)	(8.8)
(2) Replace EU	C × AI Growth	−0.8	−2.5	6.0	−0.1
		(5.0)	(7.7)	(7.0)	(8.7)
(3) Horse race	EU × AI Growth	31.6***	39.5***	43.4***	44.9***
		(8.6)	(9.9)	(7.4)	(9.0)
	C × AI Growth	−0.4	−2.0	6.1	−0.2
		(4.9)	(7.6)	(7.0)	(8.7)
Panel B: C = EU IT/data compliance growth (excl. US/UK)					
(1) Baseline	EU × AI Growth	31.8***	39.7***	43.4***	45.2***
		(9.0)	(10.3)	(7.2)	(8.8)
(2) Replace EU	C × AI Growth	−2.1	−3.7	4.0	−3.3
		(5.7)	(8.8)	(8.0)	(9.9)
(3) Horse race	EU × AI Growth	31.6***	39.6***	43.3***	45.0***
		(8.8)	(10.2)	(7.2)	(9.0)
	C × AI Growth	−2.1	−3.5	3.8	−3.7
		(5.8)	(8.8)	(8.2)	(9.9)
Industry FE (FF48)		Yes	Yes	Yes	Yes
Firm controls		Yes	Yes	Yes	Yes
Observations		769	769	769	769

Notes. This table reports OLS regressions of cumulative abnormal returns (CAR, in bps) around the 21 April 2021 EU AI Act proposal. C is one of two compliance hiring growth measures, z-scored in the analysis sample. *Panel A* uses the Broader measure (any core-tier compliance match, US/UK statutes included). *Panel B* uses the EU measure (core-tier matches restricted to non-US/UK-statute categories). Specification (1) is the baseline EU × AI interaction. Specification (2) replaces EU revenue share with C. Specification (3) includes both interactions simultaneously. Standard errors clustered at FF48 in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

8.1 Sample Restriction and Key Variables Construction

We restrict the sample in this test to the subsample: firms for whom both the Act binds meaningfully and the reallocation margin is economically relevant. First, we keep only HighAI firms—those whose AI hiring growth over 2015–2019 lies above the sample median—because they are most exposed to the AI Act and the related requirement. Second, within HighAI sample we focus on HighEU firms: those whose 2015–2019 average EU revenue share exceeds the sample median. Only a firm with its own deep EU footprint can expand its EU share within its revenue. The intersection HighEU∩HighAI contains 156 firms in the analysis sample.

Defining Peers. For each firm i in this HighEU $\cap$ HighAI subsample, we identify peers via a within-industry leave-one-out design. We assign i to the finest SIC partition that yields at least five total industry peers (firms of any EU or AI type): 4-digit SIC first, then 3-digit, then 2-digit. If no SIC level yields five peers, we fall back unconditionally to i's industry. Among all industry peers in the assigned peer set $\mathcal{P}_i$, we then count the HighAI firms and compute the fraction of those that are LowEU (EU share below the sample median). We define the continuous variable

$$\text{LowEUPeer}_i \;=\; \frac{\#\{p \in \mathcal{P}_i : p \text{ is LowEU and HighAI}\}}{\#\{p \in \mathcal{P}_i : p \text{ is HighAI}\}},$$

where $\mathcal{P}_i$ denotes i's peer set. The measure is defined for the 634 of $N = 771$ firms whose peer set contains at least three HighAI firms; for the remainder the HighAI denominator is too small to yield a reliable fraction. The variable lies in $[0, 1]$: higher values mean i's AI-Act-affected peer fringe is dominated by LowEU competitors. All HighEU, HighAI, and LowEUPeer constructions use only 2015-2019 data to preserve pre-determination.

Measuring EU Revenue Reallocation. The main outcome variable is the change in the within-firm EU-versus-non-EU revenue reallocation measure:

$$\begin{aligned}
\text{realloc}_{i,t} \;&=\; \log\Big(\text{EU_Rev}_{i,t}\Big) \;-\; \log\Big(\text{nonEU_Rev}_{i,t}\Big), \\
\Delta\text{realloc}_i \;&=\; \overline{\text{realloc}_{i,t\in 2021\text{–}2025}} \;-\; \overline{\text{realloc}_{i,t\in 2015\text{–}2020}}
\end{aligned}$$

where $\text{EU_Rev}_{i,t}$ is EU revenue in USD and $\text{nonEU_Rev}_{i,t} = \text{total_sale}_{i,t} - \text{EU_Rev}_{i,t}$ is Compustat sales net of EU revenue, for firm i in year t. The log-difference form compares EU markets to non-EU markets directly, and the transform nets out firm-wide growth that scales both EU and non-EU revenue proportionally.[14]

Equivalently, $\text{realloc}_{i,t} = \log\Big(\text{EU_Share}_{i,t}/(1-\text{EU_Share}_{i,t})\Big) = \text{logit}\Big(\text{EU_Share}_{i,t}\Big)$, the logit of the EU revenue share. For small EU shares this is numerically close to $\log(\text{EU_Share}_{i,t})$. We use log/logit transforms rather than EU Share in levels because the share variable in levels is mechanically bounded at zero and have many large-share outliers.

[14] For the HighEU $\cap$ HighAI subsample, the median EU share is 13%, compared with the full sample of 8% (Table A3). Removing year 2020 (2015-2019 instead of 2015-2020 as pre period reallocation measure) does not change the results qualitatively.

8.2 Empirical Specification and Results

On the HighEU ∩ HighAI subsample we estimate

$$\Delta\text{realloc}_i \;=\; \alpha + \beta \cdot \text{LowEUPeer}_i + \boldsymbol{\gamma}' \mathbf{X}_i + \delta_s + \varepsilon_i, \tag{2}$$

where $\mathbf{X}_i$ contains the same firm controls as in Equation (1), and δ_s denotes Fama–French 48 industry fixed effects.

Table 13 reports the baseline results (column 1) alongside two alternative-outcome robustness columns. Both the reallocation outcome (column 1) and the change in log EU share (column 2) show the relative within-firm shift in revenue toward EU markets versus non-EU markets. The near-equivalence of the two point estimates (around +0.4) reflects the fact that for firms with small EU shares, i.e., the median 13.2% EU revenue share, the logit and log transformations are mathematically close approximations of one another. The consistency across these two specifications strengthens the interpretation that firms facing weaker EU-embedded peers systematically reallocate revenue toward the EU post-announcement, conditional on being HighAI and HighEU themselves. By contrast, the change in log EU revenue (column 3) is far noisier and statistically uninformative. Raw EU revenue growth reflects not only EU-versus-non-EU reallocation but also firm-wide scale changes (e.g., organic growth, acquisitions) that affect both EU and non-EU segments proportionally and are orthogonal to the peer-displacement channel. It confirms that the relevant margin is the within-firm EU-versus-non-EU reallocation, not absolute EU revenue levels.

Table 13: LowEUPeer Effect On Within-Firm EU Reallocation

	(1) Δrealloc	(2) Δ log(EU_Share)	(3) Δ log(EU_Rev)
LowEUPeer	+0.428**	+0.386*	+0.210
(*t*)	(1.80)	(1.76)	(0.58)
Wild-bs *p*	0.047	0.057	0.586
Industry FE	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes
Observations	156	156	156

Notes. This table reports OLS regression of Equation 2. Column (1) reports the baseline results, and columns (2) and (3) replace the outcome with Δ log(EU_Share) and Δ log(EU_Rev), respectively. Cluster-robust *t*-statistics on FF48 industry in parentheses; wild-cluster bootstrap *p*-values (9,999 Rademacher replications) below. Stars follow the wild-bootstrap *p*. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Magnitude. The baseline estimate +0.43 log-points means that for moving from a firm whose AI-exposed peers are relatively EU-embedded (median level) to one whose peers are relatively least EU-embedded (zero) maps into about a 21% (= 0.428/2) larger expansion of EU revenue relative to non-EU revenue, or an absolute EU revenue share gain of about 2.8 percentage points ($= e^{(0.386*0.5+\ln(13.2\%))} - 13.2\%$) for the median HighEU ∩ HighAI firm with 13.2% pre-event EU share.

8.3 Interpretation and Placebo Tests

The baseline result requires HighAI and HighEU satisfied simultaneously: among HighAI firms, if the firm itself is HighEU, then peer weakness (largely LowEU peers) is associated with significant within-firm EU market-share gains for the focal firm. The double requirements point to a testable placebo.

If the firm is not in the HighEU ∩ HighAI sample, either not EU-embedded (LowEU ∩ HighAI), or not meaningfully exposed to the AI Act (LowAI, regardless of own EU share), then market share gain doesn't materialize: a LowEU firm in the same AI sector lacks the EU jurisdictional capital required to increase its within-firm EU market share, and the Act-driven displacement force is not operative for LowAI firm. The prediction: a null effect of LowEUPeer. Table 14 reports the LowEUPeer coefficient from Equation 2 estimated on the pooled non HighEU ∩ HighAI sample. The placebo coefficient is essentially zero and not significant, consistent with the double requirement of HighEU and HighAI for the market share gain channel to work.

Table 14: Placebo Tests: LowEUPeer Effect On Within-Firm EU Reallocation

Sample	Observations	LowEUPeer coef (t)	Wild-bs p
HighEU ∩ HighAI	156	+0.428** (1.80)	0.047
Non HighEU ∩ HighAI	443	−0.001 (0.00)	0.996

Notes. This table reports the placebo test of LowEUPeer effect on ΔRealloc using different subsamples. Each row reports Equation 2 estimated within the indicated subsample. Row 1 reproduces the baseline from Table 13, column (1). Row 2 reports the pooled non HighEU ∩ HighAI subsample as a combined falsification. FF48 industry fixed effects are included. Cluster-robust t-statistics on FF48 industry in parentheses. Stars follow the wild-bootstrap p. *, **, *** denote significance at the 10%, 5%, and 1% levels, respectively.

Though being descriptive evidence, this cross-sectional reallocation result is consistent with EU jurisdictional capital paying off through within-firm market-share reallocation. We view the evidence as suggestive of a real-economy margin of the jurisdictional capital effect.

9 Implications

The findings have broader implications for investors, regulators, and firms. For investors, the EU AI Act reveals that forward-looking markets price jurisdictional capital as an intangible asset. Market value firm's ability to operate within a regulatory jurisdiction and adapt when that jurisdiction changes the rules. In the context of AI regulation, prior EU embeddedness helps firms to navigate the evolving new requirements.

For regulators, the results highlight a design margin rather than an argument for weaker rules. Streamlined requirements could make compliance capability more broadly accessible. In fact, recent updates to the Act (June 2026) include measures to simplify and delay certain regulatory requirements, suggesting that the identified market reactions to proposal announcement may have evolved.[15]

For firms, the results are consistent with the idea that long standing presence in a regulated jurisdiction can become valuable when new rules arrive, because it has accumulated experience with the legal, institutional, and operational environment. This benefit needs not belong only to the largest firms. Even smaller firms may benefit if they have invested in durable operating depth in the jurisdiction where the rules are written and enforced. Whether this market presence translates into a institutional capital that persists across various regulatory developments, firm adaptation, or changing market conditions would require evidence across multiple regulatory episodes and over time, which we see merits future work.

Lastly, our results draw a sharp contrast with the GDPR setting, where regulatory effect is often viewed as accruing to scale. Under the AI Act, the relevant determinant is jurisdictional capital: deep operating experience in the regulated market. Although GDPR and the AI Act are both recent EU technology regulations, our findings suggest that the effects of one should not be mechanically extrapolated to the other. The GDPR lesson is largely about size-associated effects, while the AI Act lesson is about jurisdictional embeddedness and the ability to navigate a new, evolving regulatory regime.

10 Conclusion

The EU AI Act is the first comprehensive AI regulatory framework and offers a natural setting for the impact of AI regulation on firm valuation. Using predetermined exposure measures across about 800 US-listed firms around the April 21, 2021 Commission proposal, we document that firms with established EU market presence capture a positive announce-

[15] https://www.consilium.europa.eu/en/press/press-releases/2026/06/29/artificial-intelligence-council-gives-final-green-light-to-simplify-and-streamline-rules/

ment return on the EU×AI interaction of approximately +40 bps.

We interpret this as the priced value of jurisdictional capital: the accumulated experience, routines, and interpretive knowledge that firms accumulate through sustained EU market participation, which help firm translate the Act's mandates into compliance practice. Three lines of evidence support this interpretation. The effect operates through EU revenue share rather than firm size or international diversification. Direct measures of political access do not absorb the effect or themselves predict abnormal returns, showing that lobbying is unlikely to be a driver. The effect is amplified by high-risk AI capacity buildup, stable, geographically concentrated EU presence, and prior compliance experience with EU regulations.

Beyond stock returns, we find suggestive evidence that jurisdictional capital is reflected in real-economy outcomes. Among firms that are both high AI hiring and high EU-embedded, those facing peers with weaker EU presence subsequently increase their EU revenue share, following the Act's proposal. This within-firm reallocation is consistent with EU-embedded firms leaning further into markets where their jurisdictional capital is most valuable, particularly when competing against less EU-prepared peers, suggesting that jurisdictional capital shapes real activity, not just valuations. Our results show that capital market prices jurisdictional capital as an intangible assets, and they have important implications for both the regulators and firms regarding how AI regulation reshapes firm value.

References

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics*, 151: 103745. (cite on p. 4).

Bartov, Eli, Dan Givoly, and Carla Hayn. 2002. "The rewards to meeting or beating earnings expectations." *Journal of Accounting and Economics*, 33(2): 173–204. (cite on p. 51).

Bernard, Victor L, and Jacob K Thomas. 1989. "Post-earnings-announcement drift: delayed price response or risk premium?" *Journal of Accounting Research*, 27: 1–36. (cite on p. 51).

Biguri, Kizkitza, and Jörg R. Stahl. 2025. "Who Pays a Visit to Brussels? Firm Value Effects of Cross-Border Political Access to European Commissioners." *Journal of Financial and Quantitative Analysis*, 60(2): 948–973. (cite on p. 39).

Canayaz, Mehmet, and Zhe Wang. 2026. "Crafting an AI Compass: The Influence of Global AI Standards on Firms." *SSRN Working Paper*. SSRN 4481608. (cite on p. 4).

Christensen, Hans B., Luzi Hail, and Christian Leuz. 2016. "Capital-Market Effects of Securities Regulation: Prior Conditions, Implementation, and Enforcement." *Review of Financial Studies*, 29(11): 2885–2924. (cite on p. 3).

Comunale, Mariarosaria, and Andrea Manera. 2024. "The Economic Impacts and the Regulation of AI: A Review of the Academic Literature and Policy Actions." *IMF Working Papers*, 2024(065): 1. (cite on p. 6).

Coppola, Antonio, Matteo Maggiori, Brent Neiman, and Jesse Schreger. 2021. "Redrawing the Map of Global Capital Flows: The Role of Cross-Border Financing and Tax Havens." *Quarterly Journal of Economics*, 136(3): 1499–1556. (cite on p. 42).

Dechow, Patricia M, Richard G Sloan, and Mark T Soliman. 2004. "Implied equity duration: A new measure of equity risk." *Review of Accounting Studies*, 9(2): 197–228. (cite on p. 57).

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang. 2023. "Generative AI and Firm Values." National Bureau of Economic Research NBER Working Paper 31222. (cite on p. 4).

Faccio, Mara. 2006. "Politically Connected Firms." *American Economic Review*, 96(1): 369–386. (cite on p. 3).

Fama, Eugene F., and Kenneth R. French. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, 33(1): 3–56. (cite on p. 10).

Fama, Eugene F., and Kenneth R. French. 2015. "A Five-Factor Asset Pricing Model." *Journal of Financial Economics*, 116(1): 1–22. (cite on p. 55).

Fisman, Raymond. 2001. "Estimating the Value of Political Connections." *American Economic Review*, 91(4): 1095–1102. (cite on p. 3).

Goldberg, Samuel G., Garrett A. Johnson, and Scott K. Shriver. 2024. "Regulating Privacy Online: An Economic Evaluation of the GDPR." *American Economic Journal: Economic Policy*, 16(1): 325–358. (cite on p. 3).

Johnson, Garrett A., Scott K. Shriver, and Samuel G. Goldberg. 2023. "Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR." *Management Science*, 69(10): 5695–5721. (cite on p. 3).

Klapper, Leora, Luc Laeven, and Raghuram Rajan. 2006. "Entry Regulation as a Barrier to Entrepreneurship." *Journal of Financial Economics*, 82(3): 591–629. (cite on p. 3).

Klein, Daniel, Christopher A. Ludwig, and Christoph Spengel. 2022. "Taxing the Digital Economy: Investor Reaction to the European Commission's Digital Tax Proposals." *National Tax Journal*, 75(1): 61–92. (cite on p. 4).

Livnat, Joshua, and Richard R Mendenhall. 2006. "Comparing the post-earnings announcement drift for surprises." *Journal of Accounting Research*, 44(1): 177–205. (cite on p. 51).

Peukert, Christian, Stefan Bechtold, Michail Batikas, and Tobias Kretschmer. 2022. "Regulatory Spillovers and Data Governance: Evidence from the GDPR." *Marketing Science*, 41(4): 746–768. (cite on p. 3).

Smuha, Nathalie A. 2021. "From a 'Race to AI' to a 'Race to AI Regulation': Regulatory Competition for Artificial Intelligence." *Law, Innovation and Technology*, 13(1): 57–84. (cite on p. 66).

Stigler, George J. 1971. "The Theory of Economic Regulation." *Bell Journal of Economics and Management Science*, 2(1): 3–21. (cite on p. 3).

Voigt, Paul, and Axel von dem Bussche. 2017. *The EU General Data Protection Regulation (GDPR): A Practical Guide.* Springer. (cite on p. 66).

Zeume, Stefan. 2017. "Bribes and Firm Value." *Review of Financial Studies*, 30(5): 1457–1489. (cite on p. 4).

A Data Appendix

This appendix documents the data sources, variable construction, and sample-selection procedures summarized in Section 3.

A.1 Data Sources

Job postings. The primary measure of AI exposure is constructed from Lightcast database of job postings covering 8,926 US-listed firms from 2010 to 2026. For each firm-month, we count total job postings and AI-related job postings if the skills in the posting description require AI ethics and governance postings, machine learning, natural language processing, and computer vision. Firms are identified by Compustat `gvkey`, which permits direct linkage to financial data.

Geographic revenue. EU revenue exposure is sourced from FactSet's Geographic Revenue (GeoRev) database, which reports the percentage of total revenue derived from the European Union for 10,382 US-listed firms. Our primary exposure measure is the mean EU revenue share over FY2015–2019, a fully pre-COVID window. Firms are not required to disclose in all five years: a firm with disclosure in $k \leq 5$ years enters with the mean of its k disclosed annual shares.

Stock returns. Daily stock prices and adjustment factors are from the CRSP/Compustat Merged (CCM) database. Daily total returns are computed as

$$r_{it} = \frac{\text{prccd}_{it} \times \text{trfd}_{it}/\text{ajexdi}_{it}}{\text{prccd}_{i,t-1} \times \text{trfd}_{i,t-1}/\text{ajexdi}_{i,t-1}} - 1,$$

where `trfd` is the total return factor for dividends. For firms for which CRSP does not compute `trfd` (approximately 2,200 non-dividend-paying stocks), we set `trfd` $= 1$ so that price returns equal total returns. We retain only primary securities (`LINKPRIM` $\in$ {'P', 'C'}, `LINKTYPE` starting with 'L', `iid` $= 1$). All 771 firms have continuous CRSP coverage through June 17, 2021 (day +40 relative to April 21, 2021); no firm is silently truncated from CAR calculations due to mid-window delisting.

Fama–French factors. Daily Fama–French three-factor and five-factor returns are obtained from Kenneth French's data library.

Compustat Annual. Firm-level financial controls are drawn from Compustat Annual, averaged over FY2015–2019, applying standard filters (`indfmt` = 'INDL', `datafmt` = 'STD', `consol` = 'C').

Foreign-income diversification control. Foreign income share (`pifo`/`pi`) is constructed from Compustat to distinguish EU-specific regulatory effects from general international diversification benefits. The ratio is computed only for firm-years with $|\texttt{pi}| > 0.01$ to avoid division artifacts from near-zero pretax income.

Political access. Several measures of corporate political engagement with EU institutions and member states are constructed. First, following Biguri and Stahl (2025), we download all meetings between corporate representatives and European Commissioners from the Integrity Watch platform (https://integritywatch.eu), which aggregates mandatory meeting disclosures across three Commission periods (Juncker 2014–2019, von der Leyen I 2019–2024, von der Leyen II 2024–present; 78,029 total meetings, 25,258 involving companies). We match 73sample firms to Commissioner meetings during 2015–2019. Second, we download the full EU Transparency Register (https://ec.europa.eu/transparencyregister; 17,061 organizations)—the mandatory registry for organizations engaging with EU institutions—and match 74sample firms with registration dates before April 21, 2021.

For AI White Paper consultation, we construct an indicator for submitted comments in AI White Paper consultation stage (17 sample firms). In February 2020, the European Commission published a White Paper on Artificial Intelligence and solicited public comments—the consultation that directly informed the April 2021 proposal. We download all 1,216 submissions from the Commission's "Have Your Say" portal and match submitting organizations to our sample by firm name. The consultation ran from February 19 to June 14, 2020, and received 1,216 submissions: 150 from companies, 117 from NGOs, 112 from academic institutions, 111 from business associations, and the remainder from individuals and other organizations. Data downloaded via the Better Regulation API at here. Notable US company submissions include Google, IBM, Microsoft, Intel, Adobe, Cisco, Salesforce, and Qualcomm.

For member-state-level lobby, we check if the firm appears pre-event in any mandatory member-state political-engagement channels: France, Ireland, and Germany (86firms). France's HATVP répertoire des représentants d'intérêts (mandatory since July 2017), Ireland's SIPO Register of Lobbying (mandatory since September 2015), Germany's §25(3) PartG corporate-donations register (mandatory since July 2002), or pre-event membership in BITKOM, the German trade association for the digital economy and the canonical Berlin lobby vehicle for tech firms. Collectively, these jurisdictions contain ≤ 7 firms with top-1 EU revenue in our sample (Table A18). Italy, Spain, Sweden, Denmark, and the Netherlands had no mandatory pre-event firm-level lobby register over our window.

For EU-level trade-association membership, we construct an indicator if the firm appears on the pre-event member roster of DigitalEurope, DOT Europe, or ITI—the three major EU-level digital-industry trade associations (45 firms). DigitalEurope, DOT Europe, and ITI are the EU-level digital-industry trade associations whose pre-event corporate member rosters are recoverable from archived web snapshots (Internet Archive Wayback Machine captures of 2020-12-04, 2021-01-24, and 2020-04-24 respectively, all strictly pre-event). For member-state-level lobby and EU-level trade-association membership, we include membership of any EU subsidiaries of the firm.

A.2 Financial Controls

Financial controls are constructed from Compustat, averaged across FY2015–2019 to match the exposure measurement period: log market capitalization ($\log(\text{csho} \times \text{prcc_f})$), book-to-market ratio ($\text{ceq}/(\text{csho} \times \text{prcc_f})$), leverage ($(\text{dltt} + \text{dlc})/\text{at}$), R&D intensity ($\text{xrd}/\text{sale}$), and profitability ($\text{oiadp}/\text{at}$). Book-to-market, leverage, R&D intensity, and profitability

are winsorized at the 1st and 99th percentiles within the analysis sample (i.e., after all selection filters). Log market capitalization is not winsorized, since the log transform already compresses the right-skewed distribution. Missing values of long-term debt (`dltt`), short-term debt (`dlc`), and R&D expense (`xrd`) are treated as zero; firm-years with non-positive sales are set to missing for R&D intensity; strictly positive market capitalization is required. For FF48 industry fixed effects, each firm is classified by the SIC code from its most recent Compustat record; firms without a valid match are assigned to the residual "Other" category (FF48 industry 48; 7 firms, 0.9% of the sample).

A.3 Sample-Construction Waterfall

Our sample construction is an inner-join waterfall across five filters, summarized in Table A1. We match FactSet GeoRev firms to CRSP/Compustat using CUSIP and ISIN identifiers in a waterfall merge (ISIN-based CUSIP6, then FactSet CUSIP6): 5,224 of the 10,382 FactSet firms match; the 5,158 unmatched firms are predominantly micro-cap and OTC-listed companies outside CRSP/Compustat coverage.[16] Of the matched firms, 1,739 have non-missing EU revenue data for at least one year in FY2015–2019—the remainder are firms that FactSet lists but for which geographic revenue segments are not disclosed. From the job postings database, we retain firms with at least 5 total postings in both CY2015 and CY2019, which defines AI hiring growth; AI share is computed at the annual level by summing all AI and total postings within each year before dividing. We then require valid Fama–French three-factor CARs around April 21, 2021 (estimation window $[-250, -11]$ with at least 120 trading days), non-missing Compustat controls averaged over FY2015–2019, and a pre-event share price (end of March 2021) of at least \$5. The resulting analysis sample contains **771 firms**.[17]

For the sample to check jurisdictional reach of the Act, we extend our main sample to include firms with no EU revenue disclosure, and the resulted sample size increases to **1,264 firms**. We define EU presence indicator as a dummy variable of a firm ever having

[16] We verify that the unmatched firms are genuinely absent from CRSP/Compustat rather than lost to identifier mismatches: of the unmatched firms, the majority have US ISINs and valid US-format CUSIPs, yet their CUSIP6 identifiers do not appear in the CRSP/Compustat Merged universe. Ticker-based matching is not used due to the high rate of false matches from ticker reuse (e.g., SPACs recycling tickers of acquired companies). When multiple FactSet firms map to the same `gvkey` via CUSIP, we retain the match from the higher-priority method (ISIN over CUSIP); if tied, both are dropped as ambiguous. The crosswalk is built from the most recent CRSP/Compustat Merged vintage using each `gvkey`'s latest record. We use CUSIP6 (issuer-level) rather than CUSIP9 because FactSet's identifier is issuer-scoped—the 3-digit issue-and-check suffix is immaterial for cross-share-class matching. CUSIP6 values that map to multiple `gvkey`s in CRSP (rare) are excluded from the matchable pool rather than arbitrarily assigned. For the rare FactSet records that disclose revenue under multiple share classes linked to the same `gvkey`, we retain the first occurrence by sort order. All 771 firms in the final sample have a stable `gvkey`–CUSIP6 mapping in Compustat across FY2015–2021, ensuring that Compustat controls and FactSet revenue shares align to the same underlying entity throughout the measurement period.

[17] We do not exclude financial (SIC 6000–6999) or utility (SIC 4900–4999) firms. The EU AI Act's scope is sector-agnostic: Articles 2 and 6 impose obligations on providers and deployers of high-risk AI systems regardless of their industry classification, and both sectors are material AI adopters—banks for credit scoring, fraud detection, and algorithmic trading; utilities for load forecasting, predictive maintenance, and grid optimization. Excluding either sector would narrow the economic interpretation of the jurisdictional capital effect without a principled basis.

Table A1: Sample Construction Waterfall: Main Sample

Filter	Remaining	Dropped
FactSet GeoRev universe	10,382	—
Intersection: CRSP/Compustat (CUSIP6 waterfall)	5,224	5,158
Intersection: EU revenue disclosed (FY2015-2019, $\geq$ 1 yr)	1,739	3,485
Intersection: AI hiring growth (CY2015 and CY2019, $\geq$ 5 postings each)	887	852
Intersection: Valid FF3 CARs ($\geq$ 120 estimation trading days)	810	77
Intersection: Compustat controls (non-missing, avg FY2015-2019)	798	12
Intersection: Pre-event share price $\geq$ \$5 (end-March 2021)	771	27
Analysis sample	771	—

Table A2: Sample Construction Waterfall: Extensive Margin Sample

Filter	Remaining	Dropped
FactSet GeoRev universe	10,382	—
Intersection: CRSP/Compustat (CUSIP6 waterfall)	5,224	5,158
Intersection: AI hiring growth (CY2015 and CY2019, $\geq$ 5 postings each)	1,510	3,714
Intersection: Valid FF3 CARs ($\geq$ 120 estimation trading days)	1,339	171
Intersection: Compustat controls (non-missing, avg FY2015-2019)	1,317	22
Intersection: Pre-event share price $\geq$ \$5 (end-March 2021)	1,264	53
Extensive-margin sample	1,264	—
EU revenue disclosed FY2015-2019, $\geq$ 1 yr (EU_Indicator = 1)	771	—
No EU revenue disclosed (EU_Indicator = 0)	493	—

non-missing values of EU revenue share in FactSet's Geographic Revenue database between 2015–2019.[18] This expands the analysis sample with EU revenue disclosure from $N = 771$ to $N = 1{,}264$. The 493 additional firms satisfy every other sample filter (Compustat controls, FF48 industry, AI hiring ≥ 5 postings, pre-event price $\geq \$5$, valid FF3 CARs with ≥ 120 estimation-window trading days) but have no EU revenue disclosed over the pre-event window. The sample construction is in Table A2, and summary statistics are in Table A4.

For the multi-event analysis, we hold the sample and exposure measures fixed at their April 2021 values. This ensures that variation in abnormal returns across events reflects differences in the information content of each legislative milestone, not changes in sample composition or exposure.

[18] Under SFAS 131 and its international counterpart IFRS 8, publicly listed firms must disclose geographic segments in which they earn more than 10% of revenue or have material interests (Coppola et al., 2021). FactSet's Geographic Revenue database augments these mandated disclosures with narrative footnote and in-text information from 10-K filings, so a firm with material EU exposure is typically captured (see Appendix A.XI of Coppola et al., 2021). Under this data-generating process, a firm with no EU revenue entry in FactSet GeoRev over 2015–2019 is most plausibly a firm whose EU revenue falls below the 10% threshold or is not materially important. That said, we cannot rule out that there are firms with positive EU revenue that FactSet does not capture. Under this measurement error, contamination of the latter kind would attenuate the estimated AI effect for no-EU-revenue firms toward the disclosing-firm value. The estimates therefore bound this contamination from above.

A.4 Summary Statistics: Main Sample

Table A3: Summary Statistics

Variable	N	Mean	SD	P25	Median	P75	P99
Panel A: Key Exposure Variables							
AI_Growth (ppts/yr)	771	1.432	2.819	0.000	0.532	2.202	12.035
EU_Share	771	0.098	0.089	0.034	0.081	0.135	0.442
Panel B: Cumulative Abnormal Returns (April 2021, bps)							
CAR[0,+1]	771	16.301	259.957	-102.854	11.428	130.436	668.342
CAR[0,+2]	771	-18.861	326.153	-164.386	-11.787	130.964	875.575
CAR[-1,+1]	771	56.369	326.296	-107.603	72.002	216.598	949.879
CAR[-2,+2]	771	10.668	432.706	-186.424	37.142	208.123	943.127
Panel C: Financial Controls							
Market Cap (USD bn)	771	21.368	57.590	1.391	4.515	14.069	249.742
Book-to-Market	771	0.384	0.303	0.169	0.309	0.535	1.342
Leverage	771	0.259	0.183	0.111	0.248	0.365	0.841
R&D Intensity	771	0.093	0.205	0.000	0.020	0.102	1.405
Profitability	771	0.063	0.120	0.032	0.076	0.119	0.307
Panel D: Diversification							
Foreign Income Share	683	0.305	2.351	0.054	0.309	0.698	3.938

Notes. Descriptive statistics for the N=771 analysis-sample firms. AI_Growth is the annualized change in the AI share of job postings from CY2015 to CY2019 (percentage points per year). EU_Share is the firm mean FactSet GeoRev EU revenue share over FY2015–2019. Market Cap (USD bn) is the antilog of the log market-cap control (firm mean of log market capitalization over FY2015–2019); the regressions use log(market cap). Book-to-market, leverage, R&D intensity, and profitability are reported after 1%/99% within-sample winsorization.

Table A4: Summary Statistics: Extensive Margin Sample

Variable	N	Mean	SD	P25	Median	P75	P99
Panel A: Key Exposure Variables							
AI_Growth (ppts/yr)	1,264	1.145	2.542	0.000	0.317	1.534	11.184
EU Share (0 if non-discloser)	1,264	0.060	0.084	0.000	0.024	0.098	0.383
EU Indicator (1=disclosed)	1,264	0.610	0.488	0.000	1.000	1.000	1.000
Panel B: Cumulative Abnormal Returns (April 2021, bps)							
CAR[0,+1]	1,264	13.961	260.543	-108.495	10.930	129.627	707.486
CAR[0,+2]	1,264	-25.554	379.583	-182.888	-24.980	123.324	894.318
CAR[-1,+1]	1,264	69.719	332.326	-91.251	83.337	227.597	957.213
CAR[-2,+2]	1,264	13.953	452.466	-184.731	29.084	207.729	999.225
Panel C: Financial Controls							
Market Cap (USD bn)	1,264	16.331	47.838	1.073	3.130	11.078	222.056
Book-to-Market	1,264	0.430	0.316	0.193	0.369	0.627	1.430
Leverage	1,264	0.264	0.199	0.094	0.241	0.381	0.866
R&D Intensity	1,264	0.187	0.866	0.000	0.000	0.054	6.938
Profitability	1,264	0.051	0.126	0.023	0.066	0.109	0.309
Panel D: Diversification							
Foreign Income Share	814	0.281	2.169	0.011	0.222	0.663	3.785

Notes. Descriptive statistics for the full sample of N=1,264 firms: the 771 main-analysis firms that disclose EU segment revenue in FactSet GeoRev (EU Indicator = 1) and the 493 additional firms that do not disclose EU segment revenue (EU Indicator = 0). Under SFAS 131, firms report geographic-segment revenues only above a materiality threshold; non-disclosure therefore implies immaterial EU exposure rather than randomly missing data. EU Share is set to zero for the 493 non-disclosers. EU Indicator equals one if the firm discloses positive EU segment revenue in at least one fiscal year during 2015–2019, zero otherwise. AI_Growth is the annualized change in the AI share of job postings, CY2015 to CY2019 (percentage points per year; Lightcast). CARs are market-model cumulative abnormal returns in basis points; day 0 is April 21, 2021 (EU AI Act proposal date). Market Cap (USD bn) is the antilog of the log-market-cap control (firm mean of log market capitalization, FY2015–2019); regressions use log(Market Cap). Book-to-market, leverage, R&D intensity, and profitability are winsorized at the 1st and 99th percentiles within the full 1,264-firm sample. Foreign Income Share is available for N=814 firms.

B Robustness Appendix

B.1 Pre-Trend Trajectory: Cumulative Coefficient

Figure A.1 complements the formal pre-trend tests in Table 2 by plotting the full cumulative trajectory of the EU×AI coefficient from day −40 through day +40.

Table A5: The Interaction Effect Across Legislative Events

	Dependent variable: CAR (bps), event window $[0, +2]$			
	Apr 2021 Proposal	Jun 2023 EP Vote	Dec 2023 Trilogue	Mar 2024 EP Final
EU × AI Growth	39.7***	7.5	6.8	-16.0
	(10.3)	(14.4)	(20.4)	(13.3)
EU Revenue Share	2.4	-7.9	-18.4	-0.9
	(13.9)	(16.8)	(13.6)	(10.1)
AI Growth	-16.7**	24.2	-5.1	-3.5
	(6.4)	(16.8)	(10.2)	(13.9)
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Obs. (Apr 2021 Proposal)	769	769	769	769
Obs. (Jun 2023 EP Vote)	767	767	767	767
Obs. (Dec 2023 Trilogue)	766	766	766	766
Obs. (Mar 2024 EP Final)	765	765	765	765

Notes. This table reports OLS regressions of cumulative abnormal returns (CAR, in bps) following Euqation (1) but with four legislative events of the EU AI Act. Standard errors (FF48-clustered) are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Each event is a separate FF3 market-model re-estimation with a pre-event $[-250, -11]$ window.

Table A6: Treatment Timing and Information Leakage

	Dependent variable: CAR (bps)			
	(1) $[-5, -3]$	(2) $[-5, -1]$	(3) $[0, +2]$	(4) $[-5, +2]$
EU × AI Growth	9.1	14.6	39.7***	54.3***
	(8.6)	(11.3)	(10.3)	(13.5)
EU Revenue Share	-3.1	17.0	2.4	19.4*
	(7.3)	(10.7)	(13.9)	(11.0)
AI Growth	-8.3	-23.8*	-16.7**	-40.5***
	(8.8)	(12.9)	(6.4)	(12.0)
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Observations	769	769	769	769
RI p-value, EU×AI t	—	—	—	0.002

Notes. This standalone timing table reports the baseline cross-sectional specification over windows that start at the April 14, 2021 leak date or the April 21, 2021 official Commission release. CARs are Fama–French three-factor abnormal returns estimated over trading days $[-250, -11]$ relative to April 21. Standard errors are clustered by FF48 industry. Column (4) reports a two-sided randomization-inference p-value for the EU×AI t-statistic from placebo dates outside a ±10-trading-day buffer around April 21, 2021. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

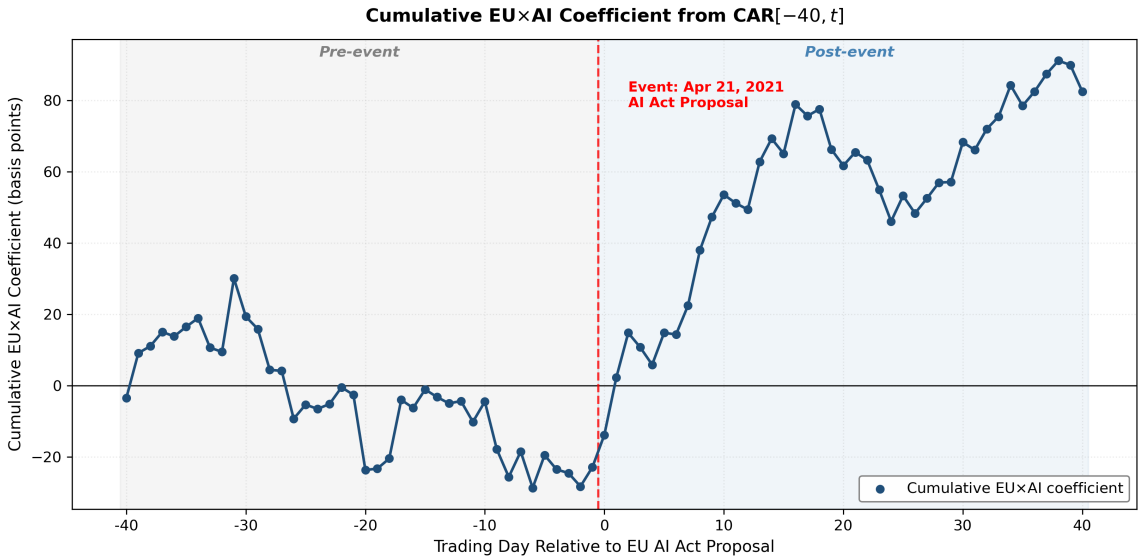


Figure A.1: Cumulative EU×AI Coefficient from CAR$[-40, t]$

Notes: Each point plots the coefficient on EU×AI (standardized) from a regression of CAR$[-40, t]$ based on Equation 1, where t ranges from day -40 to day $+40$. The cumulation anchor is fixed at day -40 throughout. The CAR estimation window $[-291, -41]$ is disjoint from the event window.

B.2 Randomization Inference: Full Distribution

Table A7 reports the full per-window results of the randomization inference procedure described in Section 5.3. Coefficient-based RI p-values for the EU×AI interaction range from 0.039 to 0.062 across the four windows. The AI Growth main effect yields t-based RI p-values between 0.091 and 0.116 across the four windows. The EU Revenue Share main effect is well within its placebo distribution on every window ($p > 0.50$). Placebo means for all three coefficients are centered near zero with no systematic bias. Figures A.2 and A.3 show the full 2×2 placebo distributions underlying these tests.

Table A7: Randomization Inference: April 21, 2021 Coefficients vs. Placebo Distribution

	$[0,+1]$	$[0,+2]$	$[-1,+1]$	$[-2,+2]$
EU × AI Growth				
Actual coeff. (bps)	+31.8	+39.7	+43.4	+45.2
Actual t-statistic	+3.54	+3.87	+6.06	+5.15
Placebo mean (bps)	+0.0	−0.0	+0.0	+0.1
Placebo SD (bps)	15.2	18.5	18.5	24.3
Placebo $[5, 95]$ (bps)	[−23.6, +23.2]	[−28.5, +28.4]	[−28.5, +28.4]	[−37.8, +38.7]
RI p-value (coeff.)	0.039	0.041	0.031	0.062
RI p-value (t-stat)	0.012	0.017	0.002	0.000
EU Revenue Share				
Actual coeff. (bps)	+3.2	+2.4	+15.1	+22.5
Actual t-statistic	+0.35	+0.18	+1.78	+2.03
Placebo mean (bps)	−0.1	−0.2	−0.2	−0.4
Placebo SD (bps)	18.5	22.6	22.6	29.9
Placebo $[5, 95]$ (bps)	[−26.5, +28.0]	[−32.5, +32.5]	[−32.6, +32.6]	[−43.7, +47.2]
RI p-value (coeff.)	0.834	0.882	0.447	0.364
RI p-value (t-stat)	0.791	0.886	0.197	0.135
AI Growth				
Actual coeff. (bps)	−21.9	−16.7	−37.1	−32.2
Actual t-statistic	−3.66	−2.60	−4.65	−4.25
Placebo mean (bps)	−0.1	−0.1	−0.2	−0.2
Placebo SD (bps)	12.6	15.2	15.2	20.2
Placebo $[5, 95]$ (bps)	[−19.6, +22.2]	[−24.6, +25.9]	[−24.4, +25.8]	[−31.2, +35.5]
RI p-value (coeff.)	0.091	0.267	0.014	0.116
RI p-value (t-stat)	0.002	0.035	0.000	0.002
Placebo dates		483		

Notes: This table compares the actual April 21, 2021 regression coefficients to a placebo distribution constructed from 483 non-event trading dates (Dec 2020 to Dec 2022, excluding ±10 trading days around April 21, 2021). For each placebo date, CARs are recomputed using a Fama–French three-factor estimation window of $[-250, -11]$ trading days relative to the placebo date, then the full interaction regression with controls and FF48 industry fixed effects is estimated. The actual coefficients and t-statistics are refreshed from the clean-room Stata `reghdfe` Table 1 output. The RI p-value (coeff.) is the fraction of placebo dates where $|\hat{\beta}_{\text{placebo}}| \geq |\hat{\beta}_{\text{actual}}|$. The RI p-value (t-stat) is the analogous fraction based on t-statistics.

B.3 Long-run Effect

Table A8 extends the baseline regression to longer windows. The EU×AI interaction grows monotonically from about +30 bps on CAR$[0, +1]$ to reaching +175 bps on CAR$[0, +40]$—more than five folds over two calendar months—and is significant at the 1% level at every horizon. The absence of reversal provides evidence against short-lived overreaction and is consistent with a permanent repricing of the jurisdictional capital.

A potential concern about Table A8 is that the long-horizon coefficient growth reflects post-event news correlated with EU×AI rather than persistent repricing of the announcement signal. We address this through a complementary test, in the spirit of local pro-

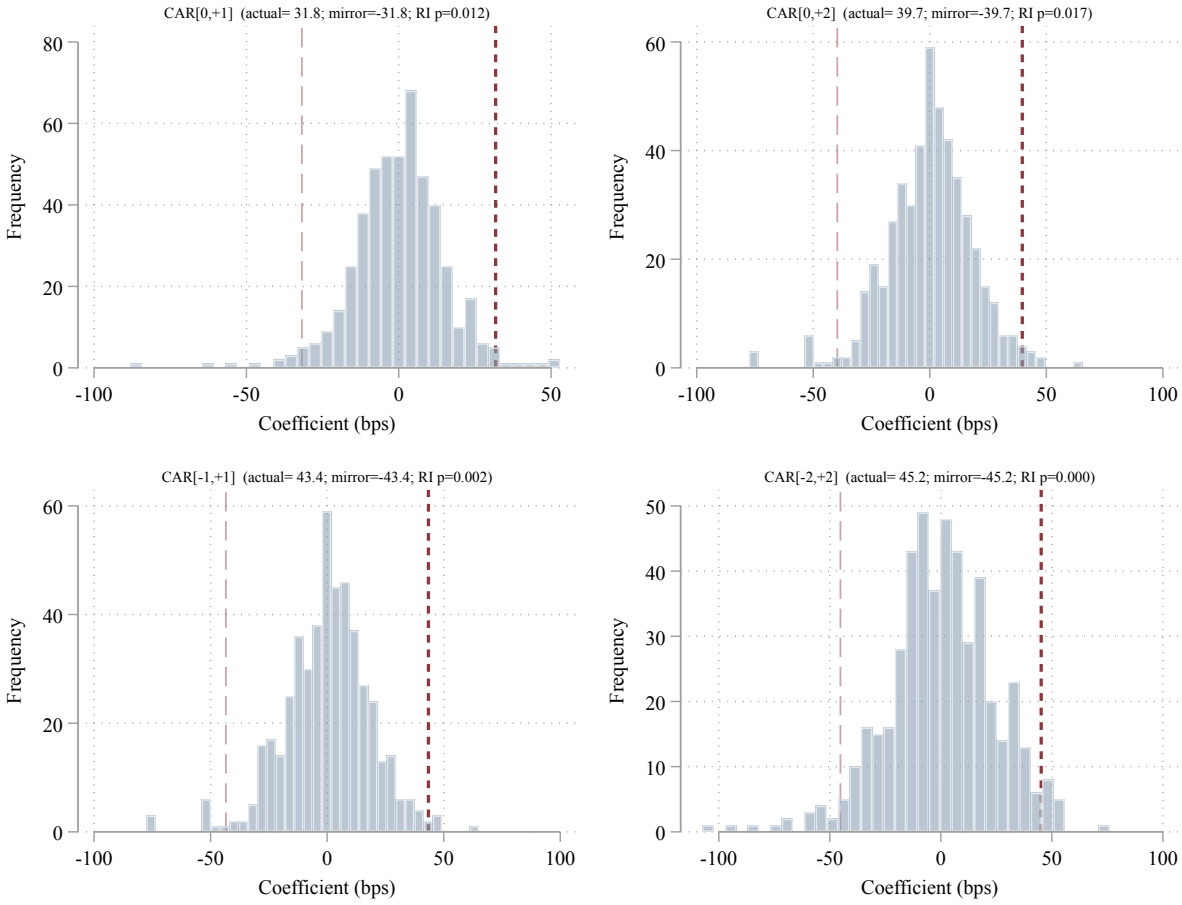


Figure A.2: Randomization Inference: EU×AI Coefficient, All Windows

Notes: Each panel plots the distribution of EU×AI coefficients from 483 placebo regressions (non-event trading dates, December 2020 to December 2022). The solid red line marks the actual April 21, 2021 coefficient; the dashed red line marks its mirror image for two-sided comparison. RI p-values are based on t-statistics.

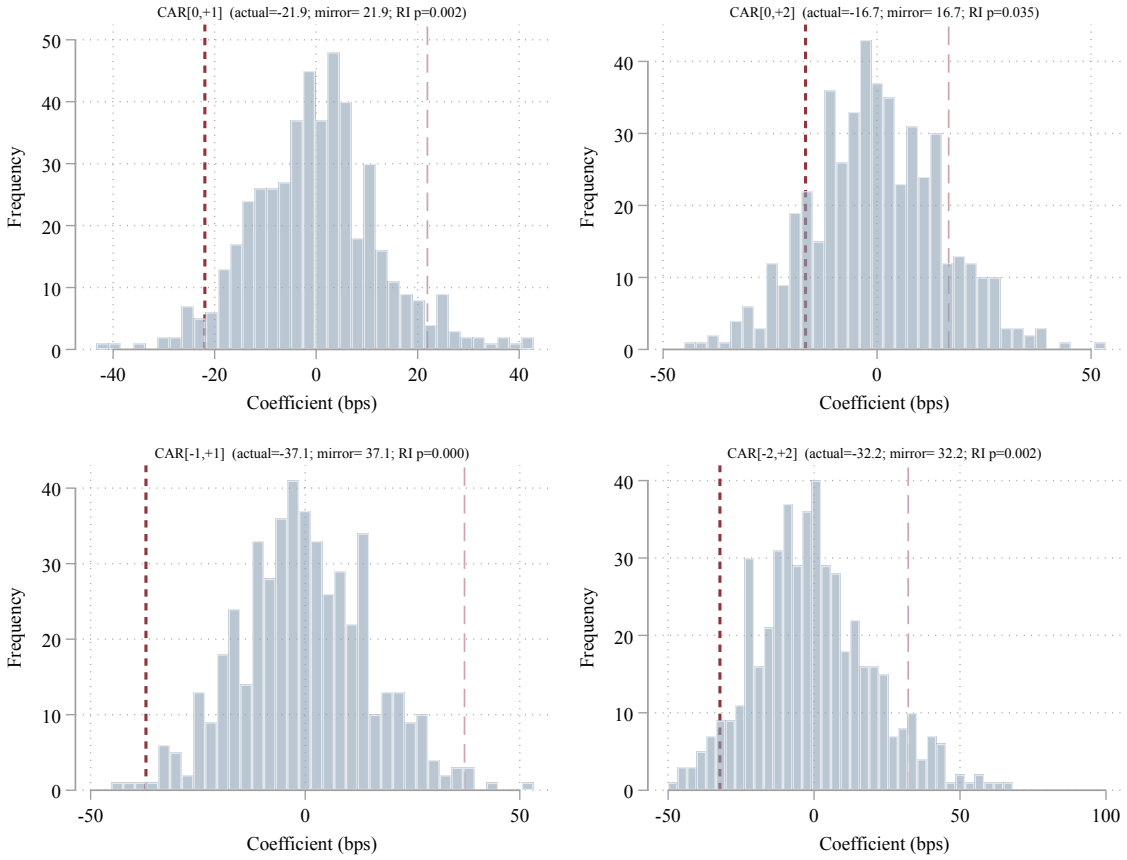


Figure A.3: Randomization Inference: AI Hiring Growth Coefficient, All Windows

Notes: Each panel plots the distribution of AI Growth coefficients from 483 placebo regressions. The solid red line marks the actual April 21, 2021 coefficient. RI p-values are based on t-statistics.

Table A8: Long-Window Persistence: EU×AI Interaction Across Event Windows

	Dependent variable: CAR (bps)						
	(1) [0, +1]	(2) [0, +2]	(3) [0, +5]	(4) [0, +10]	(5) [0, +20]	(6) [0, +30]	(7) [0, +40]
EU × AI Growth	31.8***	39.7***	43.3***	86.7***	113.0***	126.3***	174.0***
	(9.0)	(10.3)	(10.5)	(32.2)	(35.1)	(33.3)	(54.4)
EU Revenue Share	3.2	2.4	-8.6	-14.3	-5.7	17.0	16.8
	(9.4)	(13.9)	(19.0)	(26.9)	(37.1)	(28.2)	(25.2)
AI Growth	-21.9***	-16.7**	-16.4	28.4	60.9**	55.2	-1.6
	(6.0)	(6.4)	(16.6)	(25.2)	(28.1)	(33.8)	(49.5)
Industry FE (FF48)	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	769	769	769	769	769	769	769

Notes. This table reports the EU×AI coefficient from a separate OLS regression of the indicated cumulative abnormal return window, using the baseline specification Equation 1. Standard errors clustered at the FF48 industry level. *, **, *** denote significance at the 10%, 5%, 1% levels.

jection, that uses the narrowly identified announcement-window return CAR$[0, +2]$ as the regressor—thereby filtering out post-event news that affects long-horizon CAR but not the main event window—and asks whether longer-window CAR continues to drift in the same direction or mean-revert. We regress the increment CAR$[+3, +h]$ on the main-window return CAR$[0, +2]$, with baseline controls and industry fixed effects.

$$\text{CAR}_i[+3, +h] \;=\; \alpha \;+\; \theta_h \,\text{CAR}_i[0, +2] \;+\; \gamma' \mathbf{X}_i \;+\; \delta_s \;+\; \varepsilon_i, \quad h \in \{5,\ 10,\ 20,\ 30,\ 40\}, \tag{3}$$

where the LHS is the post-main-window increment return $\text{CAR}_i[+3, +h] = \text{CAR}_i[0, +h] - \text{CAR}_i[0, +2]$, the controls $\mathbf{X}_i$ and industry fixed effects δ_s are identical to the baseline specification. A positive θ_h indicates same-sign drift; $\theta_h \approx 0$ indicates the announcement news is fully impounded by day +2; a negative θ_h indicates partial reversal. Local projection tests the long-run effect from a different angle from Table A8. The cumulative EU×AI loading can rise monotonically while $\theta_h \approx 0$.

Table A9: Local Projection

CAR[+3,+h]	$h = 5$	$h = 10$	$h = 20$	$h = 30$	$h = 40$
θ	-0.0191	-0.1553**	-0.1811*	-0.0762	+0.1049
	(0.0561)	(0.0753)	(0.1054)	(0.1391)	(0.1592)
N	769	769	769	769	769
Adj R^2	0.111	0.099	0.055	0.055	0.065

Notes: Each column reports the OLS slope θ_h from Equation 3 for the indicated horizon h. Standard errors in parentheses are FF48-clustered. Significance (*, **, ***) refers to the 10%, 5%, and 1% levels.

In Table A9, θ_h is statistically indistinguishable from zero at $h = 5$, mildly negative at $h = 10$ and $h = 20$, and noisy thereafter. This suggests a temporary moderate reversal ($< 20\%$) at intermediate horizons, but no reversal once in the longer run. This shows that

the announcement-window is neither systematically reinforced nor reversed during the post-event period. Combined with Table A8, the two pieces tell a coherent story: EU×AI-treated firms continue to earn positive returns relative to controls during the post-event period, but the size of an individual firm's announcement-window move does not predict whether that firm will subsequently drift further or mean-revert.

B.4 Long-run Effect: Earnings-Announcement Contamination Robustness

The April 2021 proposal falls in Q1 earnings season. Earnings news could enter the announcement window CAR through three channels: (i) reporting status—whether the firm released earnings inside the CAR window; (ii) surprise content—size, sign, and forward-guidance implications; and (iii) post-earnings drift (PEAD) (Bernard and Thomas, 1989; Bartov, Givoly and Hayn, 2002; Livnat and Mendenhall, 2006). We address each channel in turn.

Step 1: Correlation diagnostics. We construct firm-level proxies for each channel from I/B/E/S, matched to our sample via the IBES–CRSP–gvkey link (98.8% match rate): $D_t \in \{0, 1\}$ for in-window earnings release, the standardized surprise SUE_z and its absolute value, a centered beat indicator, the 5-day post-release FY1 analyst revision, the PEAD-time exposure $\text{days_post}_{t,i} = t - \tau_i$, and signed and magnitude PEAD interactions. Table A10 reports Pearson correlations of these eight proxies with the three treatment terms (EU, AI, EU×AI) at three event horizons. The largest correlation with EU×AI is +0.11 (PEAD-time at $[0, +2]$). The strongest correlations involving the FY1 revision at the long horizon (EU: −0.26; EU×AI: +0.17) point in opposite directions and cannot jointly bias the EU×AI coefficient. Low unconditional correlations are necessary but not sufficient for unbiasedness, motivating the regression robustness ladder in Step 2.

Step 2: Regression robustness ladder. We re-estimate the baseline regression under eight nested specifications. Spec (A) is the baseline. Spec (B) adds D_t. Spec (C) adds the four earnings-news measures (with zero imputed for non-window firms). Spec (D) is the union of (B) and (C). Specs (E)–(G) progressively add PEAD-time controls. Spec (H) adds the unrestricted in-window interactions $D_t \times \text{EU}$, $D_t \times \text{AI}$, and $D_t \times \text{EU} \times \text{AI}$, identifying EU×AI from out-of-window firms only. All specs retain firm controls and FF48 industry fixed effects.

Table A11 reports the EU×AI coefficient across all eight specifications at $[0, 0]$, $[0, +1]$, and $[0, +2]$. The coefficient is invariant across A–G at every short window. At the headline $[0, +2]$ window, the range across A–G is $[+37, +40]$ bps; at $[0, +1]$, $[+30, +34]$ bps; at $[0, 0]$, $[+11, +13]$ bps. All seven restricted coefficients are significant at the 1% level. Adding PEAD-time controls in (E)–(G) does not attenuate the coefficient. The unrestricted Spec (H) coefficient is somewhat lower (range $[+10, +32]$ bps), as expected since it identifies EU×AI from out-of-window firms; it remains significant at $[0, +2]$ at the 1% level.

The Spec (H) joint F-test of $D_t \times \{\text{EU}, \text{AI}, \text{EU} \times \text{AI}\} = 0$ does not reject at the headline $[0, +2]$ window ($F(3, 44) = 1.35$, $p = 0.27$) or at $[0, 0]$ ($p = 0.37$), justifying the homogeneity

Table A10: Correlation of Earnings-Control Proxies with Treatment Terms

Earnings-control variable	EU	AI	EU × AI
Panel: window $[0, +2]$ *(*$D_t = 1$ *for 73 firms)*			
D_t (in-window indicator, 0/1)	+0.031	+0.056	+0.076*
SUE_z (signed surprise)	+0.007	-0.022	+0.019
$\lvert\text{SUE}\rvert_z$ (magnitude)	-0.040	-0.080*	-0.005
Beat (centered indicator, range $[-0.5, +0.5]$)	+0.065	+0.060	+0.099**
5-day FY1 revision_z (guidance)	+0.008	-0.042	+0.070
days_post (PEAD-time)	+0.047	+0.054	+0.114**
days_post × SUE_z (signed PEAD)	-0.002	-0.026	+0.034
days_post × $\lvert\text{SUE}\rvert_z$ (magnitude PEAD)	-0.038	-0.065	-0.010
Panel: window $[0, +10]$ *(*$D_t = 1$ *for 472 firms)*			
D_t (in-window indicator, 0/1)	+0.014	+0.028	+0.045
SUE_z (signed surprise)	-0.039	+0.032	+0.028
$\lvert\text{SUE}\rvert_z$ (magnitude)	-0.004	-0.041	-0.014
Beat (centered indicator, range $[-0.5, +0.5]$)	+0.048	+0.088*	+0.059
5-day FY1 revision_z (guidance)	-0.056	-0.001	+0.046
days_post (PEAD-time)	+0.031	+0.050	+0.092*
days_post × SUE_z (signed PEAD)	-0.025	-0.014	+0.030
days_post × $\lvert\text{SUE}\rvert_z$ (magnitude PEAD)	-0.053	-0.069	+0.008
Panel: window $[0, +40]$ *(*$D_t = 1$ *for 712 firms)*			
D_t (in-window indicator, 0/1)	+0.047	-0.012	-0.016
SUE_z (signed surprise)	-0.025	-0.015	+0.039
$\lvert\text{SUE}\rvert_z$ (magnitude)	-0.012	-0.054	+0.002
Beat (centered indicator, range $[-0.5, +0.5]$)	+0.061	+0.065	+0.005
5-day FY1 revision_z (guidance)	-0.260***	+0.001	+0.168***
days_post (PEAD-time)	+0.056	-0.007	+0.020
days_post × SUE_z (signed PEAD)	-0.026	-0.007	+0.037
days_post × $\lvert\text{SUE}\rvert_z$ (magnitude PEAD)	-0.014	-0.052	-0.001

Notes: Pearson correlations between standardized treatment terms (EU, AI, EU×AI, z-scored within sample) and earnings-control proxies, computed at three event windows. Earnings-control proxies are zero for firms outside D_t. The Beat (centered) variable is $\mathbf{1}\{\text{SUE} > 0\} - 0.5$. *, **, *** denote significance at 5%, 1%, 0.1%.

Table A11: Earnings-Contamination Robustness Ladder

Spec	Earnings controls included	$[0,0]$	$[0,+1]$	$[0,+2]$
(A)	Baseline (no earnings controls)	$+10.96^{***}$	$+31.80^{***}$	$+39.69^{***}$
		(3.77)	(8.99)	(10.25)
(B)	D_t only (in-window indicator)	$+11.45^{***}$	$+30.30^{***}$	$+38.28^{***}$
		(3.70)	(10.13)	(11.56)
(C)	Earnings-news controls only (0 imputed)[a]	$+12.45^{***}$	$+31.45^{***}$	$+36.60^{***}$
		(3.50)	(10.04)	(11.37)
(D)	D_t + earnings-news controls (parsimonious)	$+12.54^{***}$	$+32.52^{***}$	$+36.57^{***}$
		(3.48)	(10.51)	(11.26)
(E)	+ days_post (PEAD time, level)	$+11.99^{***}$	$+33.59^{***}$	$+38.01^{***}$
		(3.52)	(11.14)	(11.79)
(F)	+ days_post × SUE_z (signed PEAD)	$+11.58^{***}$	$+33.55^{***}$	$+38.05^{***}$
		(3.46)	(11.15)	(11.76)
(G)	+ days_post × $\lvert\text{SUE}\rvert_z$ (full PEAD)	$+11.26^{***}$	$+33.49^{***}$	$+38.60^{***}$
		(3.39)	(11.28)	(11.56)
(H)	+ D_t interactions (unrestricted)[c]	$+10.46^{*}$	$+21.82^{**}$	$+32.13^{***}$
		(5.57)	(9.28)	(11.43)
Homogeneity F-test (Spec H)[d]		$F = 1.07$ ($p = 0.37$)	$F = 3.15^{**}$ ($p = 0.03$)	$F = 1.35$ ($p = 0.27$)
Industry FE (FF48)		Yes	Yes	Yes
Firm controls[b]		Yes	Yes	Yes
Observations		769	769	769

Notes: OLS regressions of CAR (in bps) for the indicated window on the full interaction model augmented by progressively richer earnings-news controls. Earnings-news controls are SUE_z, $\lvert\text{SUE}\rvert_z$, a centered beat indicator, and the 5-day post-release FY1 analyst revision, all zero for firms outside D_t. Spec (H) adds in-window dummy interactions $D_t\times$ {EU, AI, EU×AI}. The F-test row reports the joint F-statistic for the three Spec (H) interactions equal zero, with FF48-clustered variance. Standard errors clustered at FF48. *, **, *** denote significance at 10%, 5%, 1%.

restriction maintained in Specs (B)–(G). The test marginally rejects at $[0, +1]$ ($p = 0.03$); inspection shows the rejection comes from $D_t \times \text{AI}$ rather than $D_t \times \text{EU} \times \text{AI}$, so the jurisdictional capital itself is not the contaminated channel. Figure A.4 extends the comparison to all horizons through $[0, +40]$ and shows that baseline and PEAD-corrected coefficients track within ~ 10 bps with fully overlapping confidence bands at every horizon.

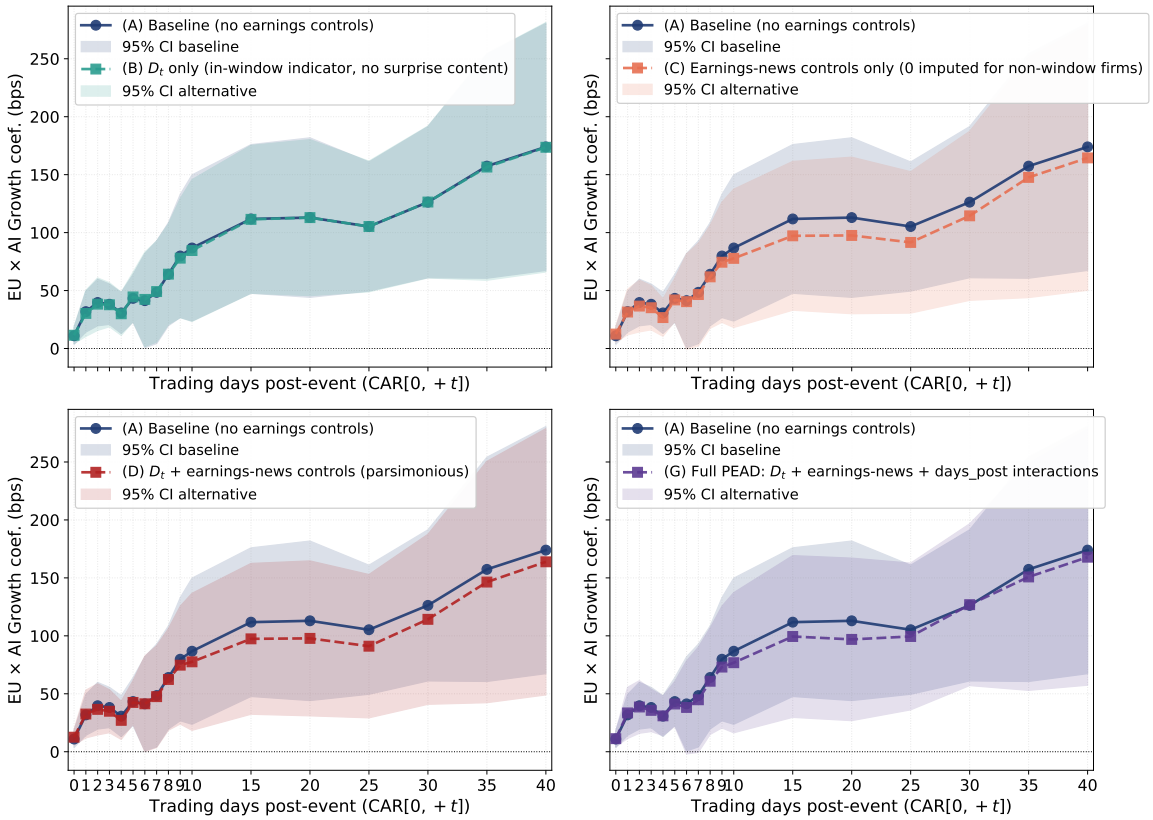


Figure A.4: EU×AI Coefficient by Horizon: Baseline vs. Earnings-Contamination Specifications

Notes: Each panel plots the baseline EU×AI coefficient (solid blue) against one contamination-control specification from Table A11 (dashed), with 95% CIs. Top-left: baseline vs. (B). Top-right: baseline vs. (C). Bottom-left: baseline vs. (D). Bottom-right: baseline vs. (G). All regressions include firm controls and FF48 industry FE with FF48-clustered SE.

Step 3: Falsification. If earnings contamination drove the headline coefficient, dropping firms whose Q1 2021 earnings release fell inside the announcement window should substantially attenuate the estimate. Excluding the 73 firms that reported during April 19–23 leaves a 698-firm subsample; the headline $[0, +2]$ coefficient is +34 bps ($t = 3.2$), a 6-bp reduction that remains significant at the 1% level. A complementary cell-level descriptive check goes further: if positive Q1 2021 surprises drove the jurisdictional capital effect, the High-EU×High-AI cell would have systematically larger surprises than the Low-EU×Low-AI control. The opposite is true—mean SUE in the jurisdictional capital cell is +0.20% of pre-event price, versus +0.33% in the Low-EU×Low-AI cell.

Bounding residual bias. The PEAD literature finds typical magnitudes of 30–60 bps over 60 trading days for a 1-SD price-deflated surprise. Applied to our 73-firm in-window subsample and the observed $|r| \leq 0.11$ correlation between PEAD-time and EU×AI, an upper bound on PEAD-driven bias to the long-horizon coefficient is roughly 20–40 bps. The observed gap between baseline and full-PEAD-corrected coefficients at $[0, +40]$ is 9 bps—within this bound and consistent with PEAD operating in the data but orthogonal to the EU×AI dimension. Non-linear contamination outside our specifications (e.g., threshold effects in PEAD correlated with EU×AI) cannot be ruled out, but the convergence of small linear correlations, regression invariance across seven specifications, drop-test stability, and a residual gap within the literature-implied bound makes such hidden contamination an unlikely explanation for the jurisdictional capital effect.

B.5 Fama–French Five-Factor Robustness

Our baseline results use the Fama–French three-factor model to estimate expected returns. To ensure that the findings are not driven by exposure to the profitability (RMW) or investment (CMA) factors, we re-estimate all abnormal returns using the Fama–French five-factor model (Fama and French, 2015) and repeat the full interaction model. Table A12 compares the three key coefficients under FF3 and FF5 across all four event windows. The results are virtually identical: the EU×AI coefficient differs by at most 1.5 bps between factor models on any window, and significance levels are unchanged. On CAR$[0, +2]$, for example, the interaction is +37.6 bps under FF3 and +39.2 bps under FF5. The AI hiring growth and EU revenue share standalone effects are equally insensitive. The stability across factor specifications confirms that the results reflect genuine cross-sectional variation in regulatory exposure, not differences in factor loadings on profitability or investment.

Table A12: Fama–French Five-Factor Robustness: Full Interaction Model

	Dependent variable: CAR (bps)			
	(1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$
EU × AI Growth	32.2***	40.8***	42.6***	44.7***
	(8.6)	(10.5)	(6.9)	(8.5)
EU Revenue Share	1.9	1.3	14.1*	21.8*
	(9.4)	(14.2)	(8.1)	(11.1)
AI Growth	-20.7***	-17.2***	-34.5***	-30.3***
	(6.0)	(6.3)	(7.3)	(6.7)
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes
Observations	769	769	769	769
Adjusted R^2	0.020	0.091	0.049	0.080

Notes. OLS regressions of cumulative abnormal returns around the 21 April 2021 EU AI Act proposal, with abnormal returns computed from the Fama–French *five-factor* market model (Mkt–RF, SMB, HML, RMW, CMA) estimated over $[-250,-11]$ trading days. All other variables and conventions match Table 3 (baseline FF3). Standard errors, clustered at the FF48 industry level, are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

B.6 Historical EU Revenue Share Robustness

Table A13: Historical EU Revenue Share

EU measure	Dependent variable: CAR (bps) (1) $[0,+1]$	(2) $[0,+2]$	(3) $[-1,+1]$	(4) $[-2,+2]$	N
Panel A					
Avg 2015-2019 (baseline)	31.8***	39.7***	43.4***	45.2***	769
	(9.0)	(10.3)	(7.2)	(8.8)	
Avg 2014-2018	33.1***	40.7***	43.8***	45.1***	758
	(9.3)	(10.3)	(7.3)	(9.5)	
Avg 2012-2016	25.0***	32.4***	36.3***	36.7***	711
	(7.2)	(8.0)	(9.2)	(9.5)	
Avg 2010-2014	11.3*	14.7**	20.2*	20.3	666
	(6.4)	(6.9)	(11.9)	(13.7)	
Avg 2008-2012	11.9*	13.8*	25.3**	22.8	620
	(6.9)	(8.0)	(11.8)	(13.9)	
Panel B					
FY2004	10.7***	11.2***	24.2**	24.7*	430
	(2.8)	(3.7)	(11.3)	(13.9)	
FY2007	5.3	1.9	21.4*	18.3	485
	(6.0)	(5.3)	(11.4)	(12.1)	
FY2012	14.8**	17.2**	30.6***	28.6**	606
	(6.9)	(8.1)	(10.7)	(12.2)	
FY2015	29.6**	35.1**	44.0***	44.3***	674
	(11.2)	(13.0)	(7.5)	(10.9)	
Industry FE (FF48)	Yes	Yes	Yes	Yes	
Firm controls	Yes	Yes	Yes	Yes	

Notes. OLS regressions of cumulative abnormal returns (CAR, in bps) around the 21 April 2021 EU AI Act proposal. Each row re-estimates the CAR$[0,+2]$ specification of Table 1 replacing EU×AI with an alternative historical EU-exposure measure interacted with AI hiring growth. Standard errors are clustered at the FF48 industry level and reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

B.7 Equity Duration and Sensitivity to News

Firms rapidly expanding AI hiring tend to be larger and more highly valued, which raises the concern that their negative announcement returns reflect long-duration growth-firm exposure to general macroeconomic news rather than AI-specific news. Book-to-market and size in the baseline partially absorb this variation, and the Fama–French model absorbs the average HML return, we further add firm-specific duration and risk-news sensitivity to control for sensitivity to general macro news.

We construct an implied equity-duration measure from fiscal-year 2019 market equity, book equity, earnings, and sales growth following Dechow, Sloan and Soliman (2004). We forecast cash flows for 15 years using the paper's standard mean-reversion parameters and a 12% discount rate. For sensitivity to general macro news, we separately estimate each firm's sensitivity (beta) to daily changes in the 10-year Treasury yield and VIX in an augmented Fama–French five-factor regression over 500 pre-event trading days ending 30 trading days before the proposal.

Table A14 reports the results. In the full FF5 sample, EU×AI is +40.8 bps. Restricting to the 722-firm common duration/beta sample changes these estimates to +33.3 bps. Adding duration and the two pre-event risk betas leaves them largely unchanged, so direct exposure controls produce essentially no attenuation. When abnormal returns themselves

are additionally purged of the pre-event yield/VIX component, EU×AI remains +30.0 bps.

Table A14: Equity Duration and News Sensitivity

	FF5 base	Common base	+ Duration	+ Risk betas	+ All	Risk-adjusted CAR
EU × AI Growth	+40.8*	+33.3*	+33.3*	+33.6*	+33.6*	+30.0*
	(10.8)	(9.4)	(9.4)	(9.4)	(9.4)	(8.7)
Wild-bootstrap p	0.086	0.067	0.067	0.064	0.065	0.062
AI Growth	-17.2*	-14.4	-14.4	-14.8	-14.8	-9.6
	(6.5)	(9.2)	(9.2)	(8.9)	(9.0)	(10.1)
Wild-bootstrap p	0.053	0.245	0.240	0.205	0.203	0.417
Duration control	No	No	Yes	No	Yes	Yes
Yield/VIX beta controls	No	No	No	Yes	Yes	No
Observations	769	722	722	722	722	722

Notes. Coefficients are basis points. Panel A uses CAR$[0,+2]$ from the Fama–French five-factor model except the final column, whose abnormal returns also partial out pre-event 10-year-yield and VIX sensitivity. Columns 2–6 use a common complete-case sample. All regressions include the standard firm controls and FF48 fixed effects. Parentheses contain FF48-clustered standard errors. Stars and reported p-values use a null-imposed Rademacher wild-cluster bootstrap with 9,999 replications.

B.8 Jurisdiction Reach

We verify the jurisdiction reach of the Act (i.e., firms with operations in EU), we separately estimate the impact of AI hiring growth on CAR on subsamples of firms with and without EU revenue disclosure.

$$\mathrm{CAR}_i[\tau_1,\tau_2] = \delta_1^{EU}\mathrm{AI}_i + \gamma'\mathbf{X}_i + \delta_s + \varepsilon_i \tag{4}$$

Table A15 Panel A reports summary statistics across the two groups. EU_Indicator= 0 firms are systematically smaller, have lower foreign-income share, are less concentrated in technology-oriented industries, and have lower AI hiring growth over the pre-period. The nearly two-fold gap in foreign-income share is direct evidence that the non-disclosing group is a predominantly domestic subset of the US-listed universe, consistent with the materiality-threshold interpretation.

The impact from the proposal announcement largely disappear for firms outside EU jurisdiction, when estimating the AI standalone effect separately within each subsample (Table A15 Panel B). Among firms with EU revenue, the AI standalone effect is negative and significant across all four CAR windows, ranging from −16 to −31 bps, matching the baseline results in Table 1. Among firms with no material EU revenue, the coefficient for AI growth is statistically indistinguishable from zero on every window, consistent with the EU revenue share as a measure of the Act's jurisdictional reach. Firms that do not sell into the EU market are not subject to the Act's conformity assessments, risk-management obligations, post-market monitoring requirements, or any other requirements, and their AI hiring growth is therefore not repriced on the April 2021 proposal.

C Post-Proposal AI-Act Lobbying

Because post-April-2021 lobbying is realized *after* the proposal event, it cannot enter the event-study regression without look-ahead bias. We therefore treat lobbying as an *outcome*

Table A15: EU AI Act Jurisdiction Reach

	Dependent variable: CAR (bps)			
	(1)	(2)	(3)	(4)
Panel A: Firm characteristics by EU_Indicator				
	$EU_{Ind} = 0$	$EU_{Ind} = 1$	$Difference$	$pvalue$
log(Market cap)	7.669	8.421	0.751***	-7.40
AI hiring growth (ppts/yr)	0.007	0.014	0.007***	-5.08
Foreign income share	0.159	0.305	0.145	-0.70
Book-to-market	0.497	0.379	-0.117***	5.70
Leverage	0.275	0.262	-0.014	1.13
R&D intensity	1.424	0.600	-0.824	1.03
Profitability	0.030	0.059	0.028***	-3.17
Tech industry (FF48 business/computers)	0.073	0.331	0.258***	-11.11
Panel B: AI Growth effect within each EU_Indicator arm				
	$[0, +1]$	$[0, +2]$	$[-1, +1]$	$[-2, +2]$
AI Growth, EU_Ind = 0 (non-disclosers)	9.8	51.7	8.9	30.4
	(23.0)	(53.4)	(24.9)	(56.8)
N	488	488	488	488
AI Growth, EU_Ind = 1 (disclosers)	-20.4***	-16.4**	-34.5***	-30.8***
	(6.6)	(6.7)	(9.1)	(8.9)
N	769	769	769	769
Industry FE (FF48)	Yes	Yes	Yes	Yes
Firm controls	Yes	Yes	Yes	Yes

Notes. Panel A compares the firms with no FactSet-GeoRev EU revenue disclosure in FY2015–2019 (EU_Ind = 0) to the disclosers (EU_Ind = 1). Panel B runs the Table 1 specification on the expanded sample with the binary EU_Indicator in place of EU revenue share and re-estimates the AI-growth-only regression separately within each EU_Indicator arm. Standard errors clustered at FF48. *, **, *** denote significance at 10%, 5%, 1%.

and test whether the April-2021 EU×AI predicts subsequent AI-Act lobbying intensity.

Outcomes. We construct four post-2021 firm-level indicators, pooling observable channels into coarse categories. *EU-level firm lobbying* (74 firms) equals one if the firm directly engaged EU institutions on the AI Act via any of: (i) formal feedback on the April 2021 proposal, (ii) a post-event AI-Act entry in the EU Transparency Register, (iii) a post-2021 Commissioner meeting on an AI subject, or (iv) signing the voluntary EU AI Pact / General-Purpose AI Code of Practice. *Country-level lobbying* (18 firms) equals one if the firm (or an EU subsidiary) appears on an AI subject in a member-state lobby register (German Lobbyregister or French HATVP). *EU-level trade-association membership* (84 firms) equals one if the firm is a member of a major EU digital-industry association that lobbied the Act (DigitalEurope, CCIA, BSA, DOT Europe, ITI, or the European Tech Alliance). *Any post-2021 AI lobbying* (115 firms) is the union of the three categories. All channels are dated after April 2021 and are used only as outcomes; we exclude non-EU-directed venues (e.g., U.S. standards bodies).

Specification and findings. We estimate linear probability models of each lobbying indicator on the standardized EU×AI interaction, the standardized EU revenue share and

AI hiring growth main effects, log market capitalization, standard firm controls, and FF48 industry fixed effects, with FF48-clustered standard errors; inference for the EU×AI term uses a wild-cluster bootstrap (9,999 replications). The EU×AI interaction does not predict post-proposal lobbying in any category: a one-standard-deviation increase shifts EU-level firm lobbying by +0.9 pp (wild-bootstrap $p = 0.44$), country-level lobbying by +1.0 pp ($p = 0.73$), trade-association membership by +0.8 pp ($p = 0.25$), and any lobbying by +0.8 pp ($p = 0.59$). Replacing EU×AI with the realized April-2021 abnormal return yields the same null pattern.

These nulls are not driven by low power: in the same regressions, AI hiring growth and firm size strongly predict lobbying (AI hiring growth: +3.3 to +4.3 pp, $p < 0.01$; log market cap: $t \in [2.9, 8.0]$), whereas the EU revenue-share main effect is also flat ($|t| \leq 1.0$). Thus AI-intensive and large firms lobby the AI Act, but AI-intensive and high-EU-embedded firms do not lobby more by virtue of the EU×AI combination. The evidence supports an interpretation of the effect as compliance/implementation capability tied to embedded EU operations, rather than political advocacy to shape the Act, with the caveat that our measures capture documented channels and may miss informal engagement or activity shifted to the post-2024 implementation phase.

Table A16: Post-Proposal AI-Act Lobbying

	(1) EU-level firm lobbying	(2) Country-level lobbying	(3) EU-level trade association	(4) Any post-2021 AI lobbying
EU × AI Growth	+0.91	+0.98	+0.80	+0.84
	(1.22)	(0.58)	(0.83)	(1.46)
EU Revenue Share	+0.45	+0.35	−0.12	−0.17
	(0.77)	(0.36)	(0.64)	(0.84)
AI Growth	+4.26***	+0.91*	+3.28***	+3.90***
	(1.20)	(0.52)	(1.10)	(0.90)
log Market Cap	+5.85***	+2.10***	+6.81***	+8.36***
	(1.49)	(0.72)	(1.03)	(1.04)
Wild-bootstrap p (EU×AI)	0.44	0.73	0.25	0.59
Controls, FF48 industry FE	Yes	Yes	Yes	Yes
Firms (lobbying = 1)	74	18	84	115
Observations	769	769	769	769

Notes. Each column is a linear probability model of the indicated post-2021 AI-Act lobbying indicator on the standardized EU×AI interaction, the standardized EU revenue share and AI hiring growth main effects, controls, and FF48 industry fixed effects. Coefficients are percentage-point changes in the lobbying probability per one-standard-deviation increase in the regressor; cluster-robust t-statistics (FF48 industry) in parentheses. Stars on the EU×AI row follow the wild-cluster bootstrap (9,999 replications; the bootstrap p is reported in the "Wild-bootstrap p" row); stars on the remaining rows follow the cluster-robust t. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D High-Risk-AI Hiring Dictionary

A Lightcast job posting is classified as *high-risk AI* if at least one listed skill matches the high-risk-AI skill dictionary *and* its job title matches the high-risk-AI title dictionary (AND-logic). Both dictionaries map the Lightcast universe to the high-risk areas of the April 2021

EU AI Act proposal (COM(2021) 206 final); the rule is dictionary-based rather than topic-model-based to ensure full replicability. Table A17 reports per-category counts and sample entries.

Included scope. Five high-risk areas of the 2021 proposal are mapped: Annex III(1) biometric identification (remote identification and verification systems); Annex III(3) education and vocational training (access determination and student assessment); Annex III(4) employment and worker management (recruitment, evaluation, and performance monitoring, subject to the HR/recruiting exclusions below); Annex III(5)(b) creditworthiness assessment (credit risk modeling and mortgage origination); and Article 6(1) Annex II Section A safety components (robotics and medical-device AI under the Machinery Regulation (2006/42/EC) and the Medical Devices Regulation (2017/745)).

Excluded scope. Three groups are omitted. First, Annex III categories subject to deployer restrictions in the 2021 proposal—public-benefit eligibility, law enforcement, migration, and administration of justice—are excluded because US-listed corporate firms rarely satisfy the deployer condition: a corporate posting does not put an AI system into use by a law-enforcement or judicial authority. Second, Annex III(2) critical infrastructure is excluded because its 2021-proposal scope is narrow and generic industrial postings fail the safety-component test. Third, Annex II Section B safety components (motor vehicles, civil aviation, rail) are excluded because Article 2(2) defers the Act's main high-risk obligations to sectoral type-approval regimes, so the 2021-proposal conformity-assessment burden does not apply.

Excluded HR/recruiting patterns. Within Annex III(4), generic HR/recruiting patterns—talent acquisition, people analytics, automated recruitment, and resume parsing—are excluded because AI co-occurrence in these postings reflects general HR-tech adoption rather than targeted high-risk-AI capability building. Technical performance management variants (network, sales, operational, enterprise, and supplier performance management; Oracle EPM; PeopleSoft EPM) and applicant tracking system skills are retained. The full set of excluded patterns is preserved in a quarantine bucket in the replication package for transparency.

Dictionary summary. After these exclusions, the headline dictionary contains 128 entries (56 skills, 72 titles) across the five Annex categories.

Table A17: High-Risk-AI Hiring Dictionary: Per-Annex-Category Counts and Sample Entries

Annex III category (2021 proposal)	Skills	Titles	Sample dictionary entries
Annex III(1) Biometric identification	10	3	*Skills:* Biometrics; Fingerprint Identification And Classification; Fingerprint Recognition; Biometric Passport; Certified Biometrics Professional. *Titles:* Biometric Health Screeners; Biometrics Technicians; Biometric Screeners.
Annex III(3) Education and vocational training	2	1	*Skills:* Educational Assessment; Admissions Operations. *Titles:* Directors of Admissions Operations.
Annex III(4) Employment, workers' management (post-HR cleanup)	12	5	*Skills (ATS):* Applicant Tracking Systems; CATS Applicant Tracking System. *Skills (technical PM):* Operational/Network/Sales/Enterprise/Supplier Performance Management; Human Performance Evaluation; Oracle EPM; PeopleSoft EPM. *Titles:* Performance Management Analysts/Directors/Consultants/Specialists/Managers.
Annex III(5)(b) Creditworthiness assessment	10	14	*Skills:* Credit Risk Modeling; Credit Risk Analysis; Credit Risk Management; Mortgage Loan Origination; Commercial Mortgage Loan Origination. *Titles:* Credit Risk Managers; Credit Risk Analytics Managers; Credit Risk Underwriters; Credit Risk Associates; Credit Risk Analysts.
Art. 6(1) + Annex II Section A safety components	22	49	*Skills:* Robotic Systems; Advanced Robotics; Industrial Robotics; Robotic Surgery; Computer Aided Diagnosis; Medical Imaging Physics. *Titles:* Biomedical Imaging Technicians; Medical Imaging Technologists; Robotics Coaches; Biomedical Imaging Specialists; Medical Device Territory Managers.
Total (headline)	**56**	**72**	**128 entries**

Notes. A posting is classified as high-risk AI if at least one skill matches the skill dictionary *and* its title matches the title dictionary (AND-logic). Counts are after the HR/recruiting cleanup described in the appendix text. Excluded Annex III categories (2 critical infrastructure, (5)(a) public-benefit eligibility, (5)(c) emergency dispatch, 6 law enforcement, 7 migration, 8 justice) and Annex II Section B safety components under the Article 2(2) carve-out are documented above.

E Country Decomposition of the HHI-Based Subsamples

A natural question about the concentration-based jurisdictional capital effect reported in Section 7.3 is whether it reflects broad EU regulatory capital (accumulated experience with EU-level rulemaking at the Commission and the European AI Board) or capital specific to a single national jurisdiction (for example, a Germany-only or Ireland-only effect that happens to load on the HHI measure). Table A18 decomposes each HHI-based subsample by the EU member country in which its firms earn revenue. The pattern is unambiguous:

both subsamples are Germany-dominant. Germany carries 33.6% of total EU revenue in the high-HHI half and 21.5% in the low-HHI half, and Germany is the top-1 EU country for 83% of high-HHI firms (320 of 386) and 98% of low-HHI firms (379 of 385). What separates the subsamples is not which countries they operate in but how tightly they concentrate within those countries: the high-HHI subsample concentrates its EU revenue more sharply in Germany, France (20.8% vs. 15.5%), and Italy (10.3% vs. 13.8%), while the low-HHI subsample spreads the same total more evenly. Ireland is slightly over-represented in the high-HHI subsample (11 firms top-1 vs. 2 firms), consistent with a small number of EU tax-domiciled technology firms, but Ireland accounts for only 4.2% of high-HHI EU revenue—too little to drive the +47 to +86 bps effect reported in Section 7.3. The concentration effect therefore operates through depth of engagement with the major continental EU jurisdictions whose national competent authorities will implement and enforce the AI Act, consistent with a regulatory-capital interpretation at the country-authority level rather than at the Brussels-lobbying level.

Table A18: EU Member Country Decomposition of the HHI-Based Subsamples

	High-HHI subsample (concentrated)			**Low-HHI subsample** (diversified)		
Country	Agg. share	Mean share	# top-1	Agg. share	Mean share	# top-1
Germany	33.6%	30.4%	320	21.5%	21.7%	379
France	20.8%	18.9%	24	15.5%	15.7%	4
Italy	10.3%	12.5%	6	13.8%	13.9%	0
Spain	8.1%	9.3%	8	9.7%	9.7%	0
Ireland	4.2%	4.2%	11	2.7%	2.7%	2
Belgium	3.1%	3.2%	2	3.8%	3.8%	0
Sweden	3.1%	3.1%	1	4.0%	4.0%	0
Poland	3.0%	3.9%	5	4.0%	4.0%	0
Austria	2.0%	2.3%	0	3.3%	3.2%	0
Denmark	1.6%	2.0%	2	2.6%	2.6%	0
Czech Republic	1.3%	1.3%	1	2.3%	2.2%	0
Romania	1.2%	1.5%	2	2.1%	2.1%	0
Other (15 countries)	7.9%	7.2%	4	14.7%	14.3%	0
Total	100.0%	–	386	100.0%	–	385

Notes. Decomposition of the EU revenue held by firms in each HHI-based subsample (median split at HHI= 0.1376, $N_{\text{High}} = 386$, $N_{\text{Low}} = 385$, analysis sample restricted to firms with strictly positive bottom-up EU revenue). *Agg. share* is the subsample's total revenue in country c divided by its total bottom-up EU revenue (weighted average of within-EU shares, weighted by each firm's total EU revenue). *Mean share* is the simple cross-firm average of each firm's within-EU share in country c. *# top-1* counts firms whose mean 2015–2019 share is highest in country c. Rows sorted by the high-HHI aggregate share, with the remaining countries collapsed into "Other."

F Compliance Hiring Pre EU AI Act: Detail and Robustness

This appendix provides supporting detail for the compliance-hiring split reported in Section 7.4: the variable construction and the measure's coverage of EU regulations, the article-level mapping from GDPR to AI Act obligations, and bounds on inference under small-cluster

conditions.

F.1 Variable Construction

We construct two firm-year compliance hiring measures from a curated regex screen over the Lightcast distinct-skills (31,542) and distinct-titles (72,194) dictionaries. Both measures count the share of a firm's postings whose skill or title list contains at least one core-tier match. The narrower *EU IT/data compliance* measure restricts to postings whose core matches are not exclusively in US-statute or UK-statute categories; the *Broader IT/data compliance* measure includes US- and UK-statute matches as well. For each firm-year in 2015–2019 we count flagged postings, divide by total firm-year postings, and require at least 5 total postings per firm-year (mirroring the AI hiring filter in Section 3). At firm level, the level is the mean over valid years and the growth is the four-year-annualized change.

The Section 7.4 headline uses the EU IT/data compliance measure. The two measures yield quantitatively similar patterns; the EU measure is marginally sharper, consistent with stripping out US- and UK-statute hiring without losing the GDPR-relevant signal.

F.2 Compliance Categories and EU Regulatory Drivers

The categories contributing to both measures cover six functional areas of EU compliance work in force during 2015–2019. We summarize each below.

Privacy and data protection. Data Protection Officers, Chief Privacy Officers, Privacy Counsels and Engineers, privacy-by-design practitioners, DPIA analysts, breach response staff, consent-management implementers, and privacy professional certifications (CIPP/CIPM/CIPT from IAPP, CDPSE from ISACA). The principal driver is the General Data Protection Regulation (applicable from 25 May 2018; entered into force 24 May 2016), together with its predecessor Directive 95/46/EC and the ePrivacy Directive.

Cybersecurity. Security operations (SIEM, SOAR, EDR, DLP), incident response, vulnerability and threat management, network/cloud/endpoint security, cryptography, and certifications including ISO 27001 and the privacy-extended ISO 27701. The principal EU drivers are GDPR Article 32 (security of processing), the NIS Directive, and the EU Cybersecurity Act.

Data governance and lifecycle. Data lineage and mapping, data catalogs and metadata management, data classification, master data management, records management and retention, e-discovery. The principal driver is GDPR Article 30 (records of processing activities) together with Article 5(1).

IT governance, risk, audit, and vendor management. IT general controls, GRC platforms, enterprise and operational risk management, internal IT audit, and vendor and third-party risk management. The principal drivers are GDPR Articles 24, 28, and 39

(controller demonstrability, controller–processor contracting and sub-processor due diligence, DPO compliance monitoring).

AI and algorithmic governance. Explainable AI, algorithmic accountability and fairness, AI bias testing, AI risk and safety management, model validation and monitoring. The principal driver is GDPR Article 22 (right against solely automated decision-making) together with Article 35(3)(a) and the transparency obligations of Articles 13(2)(f) and 14(2)(g). The EC's High-Level Expert Group Ethics Guidelines for Trustworthy AI (April 2019) and the AI White Paper (February 2020) shaped this category in the pre-event window.

Anti-money laundering and financial-crime compliance. AML programs, KYC and customer due diligence, sanctions screening, transaction monitoring. The EU driver is the AML Directive series (4MLD/5MLD/6MLD), although for US-listed firms in our sample, AML hiring is overwhelmingly driven by US and global frameworks (Bank Secrecy Act, FinCEN, OFAC, FATF) rather than AMLD specifically.

Several other EU instruments—eIDAS (electronic identification and trust services), PSD2 (payment services), MiFID II (financial markets), and Council of Europe Convention 108—enter the measure indirectly through adjacent skill patterns. The Broader compliance measure additionally captures core-tier patterns for US privacy and financial statutes (Sarbanes-Oxley, Gramm-Leach-Bliley, FCRA), US bank IT examination guidance (FFIEC), US export controls (ITAR, EAR), and UK transpositions of EU instruments and UK-specific regulation (UK GDPR, Data Protection Act 2018, PECR, UK NIS Regs 2018, ICO, FCA/PRA cyber-resilience, BS 7799, SMCR, UK Open Banking, Cyber Essentials).

Two methodological caveats bound the EU measure: non-EU cybersecurity frameworks (NIST CSF, FedRAMP, FISMA, HITRUST, CMMC; COBIT and ITIL) are categorized under cybersecurity rather than US- or UK-statute and therefore enter the EU measure; and GDPR's cross-border-transfer regime (Articles 44–50, including Standard Contractual Clauses, Binding Corporate Rules, and Schrems II) is substantively important but not captured by the Lightcast vocabulary.

F.3 GDPR-to-AI-Act Article Mapping

The substantive case for compliance carry-over rests on direct mappings between GDPR articles and AI Act articles. The functional layers that GDPR drove firms to build during 2015–2019 carry over as follows:

- Data governance and records of processing under GDPR Article 30 map to the AI Act's data-quality and data-management obligations under Article 10.
- Security of processing under GDPR Article 32 maps to the AI Act's cybersecurity obligations under Article 15.
- Demonstrability under GDPR Article 24 maps to the AI Act's quality-management-system requirement under Article 17 and post-market monitoring under Article 72.

- The DPO function under GDPR Articles 37–39 prefigures organizational structures relevant to the AI Act's risk-management and governance obligations.

Firms that staffed up for GDPR therefore have a plausibly lower marginal cost of complying with these AI Act articles than firms that did not. The horse race shows, however, that this carry-over is not the whole jurisdictional capital effect: compliance hiring on its own is not priced. The reason, we argue, is that the AI Act is related to but not identical to the prior EU regime. GDPR is a rules-based privacy regulation with a narrow, well-defined compliance perimeter (Voigt and von dem Bussche, 2017), whose obligations can be largely commoditized through standardized infrastructure and consulting playbooks. The AI Act is principles-based and risk-tiered, with materially broader scope and substantial interpretive ambiguity at adoption (Smuha, 2021): what counts as "appropriate" accuracy under Article 15, "adequate" risk management under Article 9, or "sufficient" human oversight under Article 14 was deliberately left to regulatory interpretation. Obligations such as conformity assessments under Article 43, the producer/deployer asymmetry under Articles 16 and 26, and the high-risk classification regime under Article 6 and Annex III have no direct GDPR analogue. What matters for these features is not only the firm's stock of explicit compliance staff but its capacity to anticipate, interpret, and shape implementation-stage decisions made by national competent authorities, notified bodies, and standards organizations—country-specific regulatory know-how that is accumulated through stable EU operations and captured by EU revenue share rather than by hiring data. Neither input is sufficient on its own: explicit compliance has nothing to leverage without jurisdictional embedding (the horse-race null), and jurisdictional embedding without operational compliance scaling delivers a smaller effect than the joint presence of both (the split's bottom half).

F.4 Compliance Category Taxonomy

Table A19 provides the example mapping of each category to its operational scope, principal EU regulation(s), and specific Article-level provisions.

Table A19: Compliance Category Taxonomy and EU Regulatory Drivers: Examples

Functional area	**Example Lightcast skills/titles**	**Principal EU regulation**	**Key provisions**
Privacy and data protection	Data Protection Officer; Privacy Engineer; DPIA analyst; CIPP/CIPM/CIPT certifications	GDPR	Articles 37–39 (DPO function); Article 25 (privacy by design); Article 35 (DPIA); Articles 12–22 (data subject rights); Articles 33–34 (breach notification)
Cybersecurity	SIEM/SOAR/EDR/DLP operations; incident response; ISO 27001/27701; CISSP, CISA, CISM	GDPR; NIS Directive; EU Cybersecurity Act	Article 32 (security of processing); Article 5(1)(f) (integrity and confidentiality); Articles 33–34 (breach notification)
Data governance and lifecycle	Data lineage and catalogs; metadata management; records retention; e-discovery	GDPR	Article 30 (records of processing); Article 5(1)(b)–(e) (purpose limitation, data minimization, accuracy, storage limitation)
IT governance, risk, audit, and vendor management	IT general controls; GRC platforms; internal IT audit; third-party risk management	GDPR	Article 24 (demonstrability); Article 28 (processor contracting and sub-processor due diligence); Article 39 (DPO compliance monitoring)
AI and algorithmic governance	Explainable AI; algorithmic accountability; AI bias testing; model validation	GDPR; HLEG Ethics Guidelines (2019); AI White Paper (2020)	Article 22 (automated decision-making safeguards); Article 35(3)(a) (mandatory DPIA for automated profiling); Articles 13(2)(f), 14(2)(g) (transparency on automated-decision logic)
AML and financial-crime compliance	KYC and customer due diligence; sanctions screening; transaction monitoring; CAMS, CAFP	AML Directive series (4MLD - 6MLD)	For US-listed firms, AML hiring is overwhelmingly driven by US frameworks (Bank Secrecy Act, OFAC); GDPR Article 6(1)(c) provides lawful basis for AML processing

Notes. Functional areas captured by the core-tier patterns in the merged Lightcast taxonomy (77 patterns across 11 categories). The example skills/titles column lists representative entries.



PUBLICATIONS

Jurisdictional Capital and AI Regulation: Evidence from the EU AI Act

Working Paper No. WP/2026/180