

AI Regulation, Startup Financing, and Talent Allocation

Junida Mulla

Saïd Business School, University of Oxford

December 2025¹

Abstract

Governments are rapidly adopting artificial intelligence (AI) regulations to balance societal risks against potential productivity gains. Using a difference-in-differences design that exploits the staggered introduction of AI-related bills across 28 U.S. states, I find that proposed regulation reduces the annual probability that an AI startup secures funding and increases the likelihood of acquisition. The effects are concentrated in rules that restrict which AI systems can be built or deployed (“constraining” rules). By contrast, documentation, auditing, and disclosure requirements (“procedural” rules) dampen these effects, consistent with standardized disclosure reducing information frictions around opaque AI technologies. Hiring patterns show that constraining rules increase demand for AI governance roles and shift technical effort from frontier AI research toward deployment of existing systems, while procedural rules have the opposite effect. Startups also reallocate AI-related jobs to non-regulated states.

Key words: AI regulation, startup financing, exit strategy, AI governance, entrepreneurship, innovation, regulatory arbitrage

¹ I would like to thank Thomas Hellmann, Ludovic Phalippou, Matthias Qian, Alan Morrison, Renée Adams, Jiaqi Zheng, Alex Montag, Joel Shapiro, Aurora Kapo, Mari Sako, Mungo Wilson, Ken Okamura, Erand Mulla as well as seminar participants at Saïd Business School for constructive feedback on this draft. An earlier version of this paper circulated under the title “AI Regulation and Entrepreneurship”.

Artificial Intelligence (AI) offers sizable potential gains in productivity and economic growth, but it also poses meaningful risks (Acemoglu, 2024; Guerreiro et al., 2023). Concerns extend beyond familiar issues such as privacy and discrimination (Comunale and Manera, 2024) to low-probability, high-impact failures, including existential risk (Acemoglu & Lensman, 2024). These worries have led some to call for a temporary freeze on frontier AI development (Future of Life Institute, 2023).² In response, regulators have stepped in with new rules and guidance. As regulatory efforts advance, however, concerns about their potential costs to innovation have intensified, prompting some critics to call instead for a freeze on AI regulation (Moens & Bradshaw, 2025; Mukherjee, 2025; Perrigo and Chow, 2025).³ Within this polarized debate, it remains unclear whether and to what extent AI regulation impedes or redirects economic progress. This paper investigates the effect of AI regulation on AI startups—key vehicles for AI innovation and diffusion (Gofman and Jin, 2023).

There are three main views about the effect of AI regulation on firms. The first view argues that regulation of AI stifles innovation. This view was influential, for example, in the debate around California’s SB 1047 proposal. The bill proposed pre-training safety clearance,

² On 22 March 2023, the Future of Life Institute published an open letter calling for “at least a six-month pause in the training of AI systems more powerful than GPT-4” and proposing that any such pause be public and verifiable, with governments stepping in to impose a moratorium if voluntary coordination failed (Future of Life Institute, 2023). Media coverage highlighted prominent signatories, including Yoshua Bengio, Elon Musk, and Steve Wozniak (Wired, 2023; Time, 2023).

³ In the United States, in May 2025, the House of Representatives included a ten-year moratorium on state and local AI regulation in its “One Big Beautiful Bill”, framing it as necessary to avoid a patchwork of state rules that could hamper AI innovation and undermine U.S. competitiveness, although the Senate later voted to remove the moratorium (Brown & O’Brien, 2025; Perrigo & Chow, 2025). In the European Union, in July 2025, the Financial Times reported an open letter by 44 European CEOs (including Airbus, ASML, SAP, and Mistral) urging the Commission to pause the EU AI Act’s application, warning that missing guidance and legal uncertainty could harm competitiveness and innovation (Moens & Bradshaw, 2025).

incident reporting, and a “full shutdown” capability for powerful AI models to circumvent catastrophic outcomes (SB 1047, California Legislature, 2023, p.10; Criddle & Hammond, 2024). Critics—including industry coalitions, major AI firms, and venture investors—argued that such mandates would front-load compliance, inject legal uncertainty, and risk pushing activity out of state, thereby chilling innovation (CCIA, 2024; OpenAI, 2024; Criddle & Hammond, 2024). Ultimately, Governor Newsom vetoed the bill, arguing that its criteria for regulating AI systems mis-targeted risk and could hinder innovation (Office of the Governor of California, 2024).

The second view, by contrast, holds that well-designed AI regulation can have a positive effect on firms in practice. By setting clear standards and disclosure obligations, regulation can narrow the information gap surrounding AI technologies, helping firms understand what is expected of them (Tartaro et al., 2023). This clarity can foster trust among users and regulators (National Institute of Standards and Technology, 2023), reduce compliance uncertainty, and support investment in AI (Tartaro et al., 2023). New York’s Advanced Artificial Intelligence Licensing Act reflects this view: it would require high-risk advanced AI systems to be registered, licensed and subject to risk-management and ethical-conduct obligations (A. 1895, New York State General Assembly, 2023). The bill promotes a “balanced” and “non-intrusive” regulatory framework as a way to foster positive innovation in advanced AI while ensuring that its benefits are realised for the state and its residents (A. 1895, New York State General Assembly, 2023, pp.1-2).

The third view is that AI regulation is largely irrelevant. The rapid pace of AI development often exceeds the government's current expertise and capacity to regulate it (Wheeler, 2023; Guha et al., 2023), leading to potentially “hollow” regulation (Guha et al., 2023, pg.5) which firms can bypass or avoid in practice. For example, firms can engage in geographic arbitrage, shifting products and operations across borders: they launch AI features first in jurisdictions with looser rules and delay or scale them back in stricter ones. Recent decisions by Meta and Apple fit this pattern, as both firms postponed the roll-out of new AI systems in the European Union due to regulatory uncertainty (Weatherbed, 2024; Campbell, 2024).

While all three views have strong proponents, what is lacking is proper empirical evidence about the effects of AI regulation on firm-level outcomes. Moreover, public debate and lobbying are dominated by large technology companies that can absorb compliance costs or shift activity across jurisdictions. For capital-constrained startups, similar regulations may be more consequential. In this paper, I therefore study whether AI regulation influences AI startups’ fundraising outcomes, exit strategies, allocation of technical effort, the geographic location of that work, and how these effects vary with regulatory design.

AI startups play a pivotal role in fostering innovation due to their rapid growth and their capacity to disrupt existing market structures through creative destruction (Gofman and Jin, 2023). Innovative young firms depend heavily on external capital to develop and commercialize new technologies (Hall and Lerner, 2010; Kerr and Nanda, 2015). Exit routes such as acquisitions offer an economic structure through which these innovations can scale and reach

broader markets (Gans et al., 2002; Gans & Stern, 2003, Phillips & Zhdanoc, 2013). Regulation that changes who gets funded or acquired therefore also changes which startups remain in the race to innovate.

A central part of the debate on AI regulation is what kind of AI gets built and where. To study this margin, I examine how regulation shapes the organization of technical work within startups. I ask whether startups continue to pursue frontier AI research or reallocate effort toward deployment and governance. I then examine where this work is carried out: whether startups maintain AI activity in the regulated state or shift AI hiring to non-regulated states. The broader question the paper addresses is thus how AI regulation affects the trajectory of innovation in the startup ecosystem. For this it considers who gets funded, who gets acquired, what kind of AI gets built, and where AI work happens.

To assess the impact of AI regulation, I analyze the staggered introduction and enactment of state-level legislation across 28 U.S. states between 2018 and 2023. Startups are particularly vulnerable to state regulations due to their limited geographic flexibility, a result of constrained resources (Lee, 2022), and the critical role of local networks during their early stages (Schutjens and Stam, 2003). My baseline analysis uses a differences-in-differences approach at the startup-year level, comparing the performance of startups in states with proposed AI regulations to those in states without such proposals. The analysis utilizes Crunchbase for startup investment data and Bryan Cave Leighton Paisner (BCLP) for information on state-level AI regulations.

I find that a proposal of AI regulation reduces the annual probability that an AI startup secures funding by about 2.3 percentage points. Relative to an unconditional funding rate of 15.6% in the sample, this corresponds to roughly a 15% decline in the likelihood of raising capital in a given year, indicating a meaningful tightening of access to finance. When proposals are enacted into law, the annual fundraising probability falls by 3.2 percentage points, or about a 20% decline relative to the same baseline. However, this effect does not extend to the total amount raised by startups that do secure funding, suggesting that regulation affects which firms get funded rather than the level of investor commitment conditional on funding. In addition, I find that regulation proposals increase the probability of a startup being acquired by 1.5 percentage points—more than twice the unconditional average acquisition rate of 0.7% in the panel—consistent with regulation shaping exit strategies and encouraging acquisition as a way to manage regulatory risk.

I next develop a taxonomy of AI regulation and test how different types of rules affect startups across states. I distinguish between procedural rules, which require documentation, assessment, auditing, and disclosure to make AI systems understandable to customers and regulators, and constraining rules, which draw hard boundaries around who is in scope, when stricter duties apply, and which practices are restricted or halted. I apply this taxonomy to the full text of each proposed AI regulation, hand-coding the presence or absence of each clause based on a close reading. Most proposals rely on procedural obligations, while constraining provisions are less common and tend to target specific high-risk uses. I then compare startup outcomes under each rule type and find that the design of regulation matters. Constraining rules are

associated with lower fundraising and a higher risk of acquisition. Procedural rules, by contrast, soften these effects, consistent with standardized disclosure reducing information frictions around otherwise opaque AI technologies. Overall, it is how AI is regulated—not just whether it is regulated—that shapes startup finance and exit.

The debate around AI regulation centers on a core tension: guarding against the risks of AI without obstructing technical progress. I study this tension by asking how startups reallocate scarce technical resources across competing objectives when regulation is proposed. Specifically, I examine how firms divide effort across three aims: managing AI risk and regulatory compliance, commercializing and deploying existing AI systems, and advancing the AI frontier. These shifts are not visible in standard outcomes such as fundraising or acquisition, so I use job postings from Lightcast as a revealed-preference measure of where the next unit of technical headcount is intended to go. I construct a new firm–year panel of AI job postings. For each posting, I classify the role as governance, applied, or frontier using a RoBERTa-based text classifier fine-tuned on a hand-labeled training set. I then link these postings to regulatory exposure at the state and year level. This allows me to track how startups rebalance their internal technical effort in response to regulation. The results show a clear directional shift. Constraining rules that impose hard limits on development and raise shutdown or enforcement risk are associated with greater emphasis on governance and applied roles and a pullback from frontier AI work. Procedural rules built around documentation and disclosure move hiring in the opposite direction.

Finally, the analysis is extended to ask not only how firms reallocate technical effort in response to AI regulation, but also where that work occurs. This question matters because differences in regulatory stringency across jurisdictions allow firms to respond to new rules not by discontinuing sensitive AI activity, but by moving that activity to places with fewer obligations. This raises the concern that “AI safe havens” may emerge, in which more permissive states attract a disproportionate share of AI development and absorb greater AI risk (Tartaro et al., 2023, p.4). To test whether regulation shifts the geography of AI work, the analysis tracks where each firm opens AI jobs each year. Around 10% of AI jobs in the data are posted in non-regulated states, and this behaviour appears for about 7% of firms. Although relocation is far from universal, the firms that do reallocate AI hiring are larger AI employers on average. Proposed AI regulation is associated with an increase in both the probability that a startup posts at least one AI job in a non-regulated state and the share of its AI hiring that occurs there. Again, this response is driven primarily by constraining rules.

Taken together, the results do not provide strong support for the view that AI regulation is irrelevant for firms. A minority of startups respond by shifting AI activity to non-regulated states, but most remain exposed to the rules in their home jurisdictions. State-level regulation is followed by systematic changes in startups’ fundraising, exit, and technical hiring, showing that these policies have real effects on how firms operate. The direction of these effects depends on how regulation is designed. Constraining rules are associated with tighter financing, a higher likelihood of acquisition, and a shift away from frontier research, in line with concerns that

regulation can slow or redirect the development of advanced AI systems. Procedural rules, by contrast, weaken these funding effects and are linked to a higher share of hiring in frontier roles, consistent with the idea that transparency and documentation requirements can ease information frictions without shutting down high-end AI work.

This paper is positive rather than normative and does not attempt to identify an “optimal” level or form of AI regulation. Instead, it maps out the concrete margins along which rules affect startups—who gets financed, who exits, what kind of work teams do, and where that work takes place. These patterns highlight real trade-offs: for example, societies have to decide whether a shift away from frontier research is an acceptable cost of reducing particular risks. Evaluating those trade-offs, and deciding which outcomes are desirable, is ultimately a political and social choice that lies beyond the scope of this paper.

Literature Review

The literature on AI regulation is still relatively limited (Comunale and Manera, 2024). Some theoretical studies focus on the balance between economic benefits and societal risks. For instance, Acemoglu and Lensman (2024) advocate for gradual technology adoption to allow for learning, noting that firms often fail to account for the societal costs of potential risks, which leads to faster adoption than is socially optimal. This under-internalization of risks supports the need for taxation and regulation to ensure a more balanced approach to deploying transformative technologies. Similarly, Guerreiro et al. (2023) evaluate regulatory strategies for AI, including Pigouvian taxes and testing protocols, aimed at mitigating risks while fostering innovation.

At the same time, the literature warns that not all regulation has the same effect on firms. Kretschmer et al. (2023, p.11) argue that broad ex-ante restrictions - rules that limit which AI systems can be developed or deployed and that require extensive compliance work before a product can launch - both raise the fixed legal costs of operating in AI and have the effect of slowing development in technically and socially “uncharted” areas, where risks are hardest to predict. These costs and slow-down pressures fall asymmetrically: large incumbents can absorb ongoing legal review, documentation, and risk management by relying on in-house compliance infrastructure, while startups typically cannot. Relatedly, Tartaro et al. (2023) emphasize that the impact of AI regulation depends on how it is designed: fragmented or poorly specified requirements can increase compliance burdens and deter investment, whereas clearer and more harmonized obligations can reduce legal uncertainty, build trust in AI systems, and make adoption more scalable.

This discussion maps directly onto the two regulatory modes studied in this paper. Constraining rules operate in the same way that Kretschmer et al. (2023) describe for broad ex-ante restrictions: they draw explicit boundaries around which AI systems can be developed or deployed, and they attach credible shutdown and enforcement risk to work that falls outside those boundaries. In line with Kretschmer’s critique, the evidence shows that such rules push startups to build internal compliance and governance capacity and to reallocate scarce technical effort away from frontier AI research and toward deploying existing systems. By contrast, procedural rules - which emphasize documentation, auditing, and disclosure rather than outright limits on what can be built - are closer to the transparency-oriented view in Tartaro et al. (2023),

in which clearer obligations and greater legibility can build trust and support adoption.

Consistent with that view, procedural rules in the data do not trigger the same defensive shift and appear mitigating in their impact on startups.

A related empirical literature studies the economic impact of digital and data governance rules, with a particular focus on privacy regulation. Canayaz et al. (2022) investigate the effects of the California Consumer Privacy Act (CCPA) on voice-AI firms. Their findings reveal that the CCPA increases the cost of trading data, which in turn reduces firms' return on assets, with smaller firms—those with a limited customer base—being disproportionately affected. A significant body of work focuses on the EU's General Data Protection Regulation (GDPR). Frey and Presidente (2024) highlight that GDPR reduces technology firms' profits across 22 industries and 31 countries, while Koski and Valamari (2020) compare EU and US firms, finding that GDPR reduces the profitability of data-intensive firms in the EU. Jia et al. (2021) demonstrate that GDPR also negatively impacts venture deals and investment for EU firms relative to US firms. These studies provide important evidence that privacy regulation can impose real costs, particularly on smaller and more data-reliant firms. However, they focus on a single, broad regulatory shock applied uniformly across an entire jurisdiction, and they evaluate primarily financial outcomes.

A broader empirical literature examines how differences in regulation across regions shape entrepreneurship. Klapper et al. (2006) show that costly entry regulations suppress new firm creation, especially in industries where entry is typically high. A parallel line of work

exploits state-level policy differences in the United States: Denes et al. (2023) study tax credits for angel investors and find effects on startup entry and performance, while Fazio et al. (2019) show that R&D tax credits influence both the quantity and quality of new entrepreneurial activity. Chen and Ewens (2025) show that the Volcker Rule's limits on bank investments in venture capital reduced funding in the most exposed U.S. regions and prompted young firms to relocate toward established venture hubs, highlighting that regulation can reshape not only who gets financed, but also where startups choose to build. More closely related, Babina et al. (2024a) use staggered adoption of Open Banking rules across countries to study how regulatory changes affect venture capital activity in fintech startups. Taken together, this work shows that regional policy differences can shape who starts firms, who gets funded, where they cluster, and which sectors attract high-growth entrepreneurship.

This paper builds on and extends the aforementioned literature in three ways. First, rather than studying a single centralized policy such as GDPR, it exploits the staggered proposal and enactment of AI regulation across U.S. states. This variation makes it possible to compare otherwise similar AI startups facing different regulatory exposure at the same point in time, which helps address common identification concerns in prior work - in particular, the difficulty of finding a credible control group and defending parallel trends when one sweeping regulation hits all firms in a region at once (Johnson, 2022).

Second, the paper moves beyond average financial performance and shows how regulation changes the path of the firm. Consistent with earlier evidence on privacy regulation,

proposed AI rules lower the probability that startups successfully raise capital. This matters because reduced access to external finance can directly limit entry and experimentation in young, high-growth sectors. But the analysis also shows that regulation shifts exit strategy: affected startups become more likely to be acquired. This points to a structural consequence of AI regulation that has not been documented in prior work - namely, pressure toward consolidation rather than independent scaling, which could concentrate AI capabilities in incumbents.

Third, the paper opens the black box of how startups adapt internally. It links specific types of AI rules to (i) hiring for governance and regulatory functions (ii) how startups allocate scarce technical headcount between frontier work and deployment activity, and (iii) where that work is located. This has two implications. First, constraining, ex-ante style rules place the heaviest operational burden on startups. Second, the challenge is not only how regulation is written but also how it is distributed across jurisdictions: when states adopt different rules, stricter states risk losing AI activity, while more permissive states potentially attract—and take on the risk of—higher-risk AI work.

This paper is structured as follows. Section I provides institutional background on state-level AI regulation in the United States and describes the data. Section II estimates the effect of AI regulation on startup fundraising and exit outcomes. Section III studies regulatory heterogeneity by distinguishing between procedural and constraining rules and testing how each affects startups. Section IV examines how startups adjust their internal technical effort in

response to regulation. Section V analyzes whether firms also respond on the geographic margin by shifting AI hiring across states. Section VI concludes.

I. Data and Context

This section begins by providing contextual background on AI regulation to frame the research question. It then proceeds with a description of the data sources and summary statistics.

A. Context on AI regulation

The rapid developments in AI have sparked ongoing discussions about regulating the societal and economic risks these technologies pose. A primary motivation for regulation is the protection of personal information. AI poses particular threats in this regard, due to the sensitivity of certain technologies to adversarial attacks and manipulation (see Chao et al., 2023; Branch et al. 2022; Zhu et al., 2024), user profiling and algorithmic discrimination (Comunale and Manera, 2024), as well as the possibility of AI inferring individual information indirectly by accessing correlated data (Acemoglu, 2021). Beyond privacy, AI regulation is increasingly motivated by broader societal and economic concerns, including competition protection, national security, and financial stability.⁴

At the federal level, the U.S. has not yet passed a comprehensive legislation targeted at regulating AI (BCLP 2024a; White and Case, 2024).⁵ However, the federal government has provided guidelines regarding AI, with the most notable example being the “Executive Order (EO) on Safe, Secure, and Trustworthy Artificial Intelligence” signed in October 2023. The

⁴ For an extensive discussion of each of these issues see Comunale and Manera (2024).

⁵ There are some existing federal laws related to AI, but with limited application such as the Federal Aviation Administration Reauthorization Act which requires a review of AI in aviation, or the National AI Initiative Act of 2020, targeted at expanding AI research and development (White and Case, 2024).

order adopts a decentralized strategy for AI regulation, assigning Agencies and Departments the responsibility to issue guidelines on particular priorities, which in turn require further legislative action to be enforceable (Comunale & Manera, 2024). Due to the slow federal response to AI regulation, a number of U.S. states have begun to propose and enact their own AI-related policies (Charkoudian et al., 2024)..

Identifying the full scope of AI regulations across U.S. states presents considerable challenges. First, there is no universally accepted definition of AI regulation, as regulatory frameworks differ significantly in how AI is defined and what areas are regulated (Comunale & Manera, 2024). This lack of consistency complicates the identification of relevant laws. Second, tracking state-level regulations is technically demanding, as each state independently publishes its legislative activities, making it difficult to systematically monitor and analyze regulatory proposals across all jurisdictions.

To overcome these challenges and enhance the rigor and replicability of my analysis, I use data from Bryan Cave Leighton Paisner (BCLP), a law firm with deep expertise in AI regulation. BCLP's comprehensive tracking methodology captures a broad range of AI-related legislative actions, focusing on regulations that directly influence the development and adoption of AI technologies in businesses, as well as consumer rights protections (BCLP, 2024a; BCLP, 2024b). Their dataset takes a wide approach to AI regulation, including laws governing automated decision-making systems, which reflects the increasing intersection between AI and automation. Importantly, BCLP excludes more narrowly focused laws—such as

quasi-regulatory actions, and sector-specific rules—ensuring the analysis is concentrated on regulations with broad applicability across industries. This targeted approach is essential for this study, as it ensures that the analysis remains focused on regulations that have a significant impact on business development and AI adoption. Additionally, since I define the timing of the shock in my analysis as the proposal of the first AI regulation, the exclusion of narrower laws helps maintain consistency in measuring broader business impacts, which are most relevant for understanding overall startup performance across multiple sectors.

BCLP’s tracking system monitors the name and status of each AI regulation in every state on a quarterly basis. By leveraging this system, I address the challenge of fragmented state-level reporting, ensuring access to a consistent and regularly updated dataset of AI-related legislations. Using the list of AI bills identified by BCLP, I access the full history of each bill via the general assembly’s website for each state, where I manually collect the precise dates of proposal, enactment, and the full text of each regulation.⁶

Insert Figure I about here

⁶ At the state level, legislation, including AI regulation, can be proposed by state legislators or sometimes by outside parties (Foster, 2023). In this paper, the proposal date is defined as the date the bill was first read/introduced in the state chambers. Where the date of first reading is not available, I use the date the bill was first referred to the relevant committee. Once read, the bill is reviewed, and may be amended (Maryland General Assembly, 2023). It then goes through a vote in the state’s chambers; if it passes, it is sent to the governor for approval or veto (Foster, 2023; Maryland General Assembly, 2023).

Figure I, Panel A presents a map showing the states that had proposed AI regulations by the end of the sample period. A total of 28 states have proposed 79 AI regulations⁷. Panel B provides a more detailed timeline of these policies, depicting the years in which regulations are proposed and enacted for each state. California was the first state to propose AI regulation in 2018, followed by a few other states in the subsequent years, with 2023 seeing the highest number of proposals.

B. AI Startups and Fundraising outcomes

For data on the formation and fundraising outcomes of startups in each state, I use Crunchbase. I define AI startups based on the methodology of Gofman and Jin (2024), who use the following keywords to filter AI start-ups: artificial intelligence, machine learning, neural networks, robotics, face recognition, image processing, computer vision, speech recognition, natural language processing, autonomous driving, autonomous vehicle, or the semantic web. To this list, I add two additional keywords: generative AI and large language model. Using these keywords, I include all AI startups founded between January 1, 2012, and December 31, 2023.⁸ I rely on Crunchbase data for identifying each startup's headquarters to determine its primary location. Startup headquarters play a central role in strategic decision-making, particularly regarding regulatory compliance (Collis et al., 2007; Chandler, 1991). Startups are often limited in their geographic flexibility due to the irreversible nature of initial location decisions, which

⁷ The analysis treats Washington D.C. in the same way as U.S. states, as it has its own set of AI regulation proposals.

⁸ I exclude startups if they fall into any of the following: 1) startups classified as not-for-profit by Crunchbase; 2) startups whose founding date is after 2012, but there are fundraising rounds before 2012, as I expect the dates to be recorded with error. I intend to extend my sample until the end of 2024 once data becomes available, so that I will be tracking startups in the 6 year window before and after the first proposed AI regulation in 2018.

involve significant investments and the development of local networks, making relocation costly and challenging in the long run (Schutjens and Stam, 2003; Lee, 2022). Therefore, the headquarters location serves as a reliable proxy for assessing a startup’s regulatory response.

Due to the recent implementation of AI regulation, the number of state-level observations is limited. Therefore, I focus on analyzing the impact of AI regulation on various startup-level performance metrics. The baseline identification strategy exploits the staggered proposals of AI regulation across the U.S. states, comparing the performance of startups in states with proposed AI regulation to those without AI regulation. I define a treated startup as any startup that was part of the sample both before and after its state proposed its first AI regulation. The control startups are the startups founded in states with no AI regulation proposals as of December 31st 2023. After applying these filtering criteria, there are a total of 13,463 AI startups in my sample: 10,957 treated startups in 28 states with AI regulation proposals, and 2,506 control startups in states without any AI regulation proposals

C. Key Variables

Table 1 provides the definitions of the variables used in the main analysis, with the summary statistics reported in Table 2.

Insert Table I about here

As a first measure of performance, I analyze two key fundraising outcomes. The first, 'Raised Eq', is a binary variable indicating whether a startup raised funding in a given year. The second, 'Amount Raised', measures the total investment amount (in USD) for years when the startup successfully secured funding. On average, 16% of startups raise funding in any given year, with an average fundraising amount of \$16 million USD for those that do.

Insert Table II about here

To assess stronger indicators of startup success, I also consider two exit outcomes: whether the startup was acquired and whether it went public through an Initial Public Offering (IPO). Acquisitions are particularly relevant as they allow smaller startups, particularly resource-limited start-ups, to bypass the high costs of regulatory compliance by being acquired by larger firms that can better manage these challenges (Ince, 2024).

To control for time-invariant startup characteristics, such as the quality of the founding team or product, I include startup fixed effects. Given that the health of the start-up ecosystem can vary across states due to differing economic conditions, I incorporate time-varying state-level controls collected from the U.S. Bureau of Economic Analysis. These controls include real GDP levels, real GDP growth rates, and disposable income per capita.

Year fixed effects are used to account for common shocks that impact fundraising outcomes for start-ups in a given year. Finally, I include start-up age fixed effects to account for the startup's maturity each year.

II. Impact of AI regulation on startup fundraising and exit

This section presents the empirical strategy and main findings on the effects of AI regulation on startup financing and exit outcomes. Section II.A details the estimation methodology. Section II.B presents the results for financing outcomes, and Section II.C reports robustness checks. Finally, Section II.D extends the analysis to exit strategies.

A. Empirical Strategy

To explore the relationship between AI regulation proposals and startup outcomes, I follow a difference-in-difference design that exploits the staggered proposal of 79 AI regulations across 28 states. I define a treated startup as any startup that was part of the sample both before and after its state proposed its first AI regulation. The control group consists of AI startups in states with no AI regulation proposals as of December 31, 2023. Equation (1) presents the baseline startup-level regression:

$$y_{ist} = \beta Proposed_{st} + Age_{ist} + \delta Z_{st-1} + \alpha_{is} + \alpha_t + \varepsilon_{it} \quad (1)$$

where y_{ist} is the outcome of startup i in state s in year t . I consider four main startup outcomes:

whether it raised funding, amount raised, and whether it has an IPO/was acquired. $Proposed_{st}$ is

a dummy variable set to 1 in the years following the proposal of the first AI regulation in state s . The model includes startup (α_{is}) and year (α_t) fixed effects. Age_{ist} are startup age dummies and Z_{st-1} account for time-varying state controls (the lagged level (logged) and growth rate of Real GDP as well as lagged log of disposable income per capita). Standard errors are clustered at the state level, to account for potential correlation of shocks affecting start-ups in the same state. In the baseline specification, I use Two Way Fixed Effects (TWFE) to estimate equation (1). The key identifying assumption is that, in absence of AI regulation proposals, there are parallel trends in the outcomes of startups in states with proposals relative to outcomes of startups in states without proposals. In robustness checks, I show similar results follow when using Sun & Abraham (2021), and Callaway and Sant'Anna (2021) for estimation.

The dynamic version of equation (1), which allows testing for parallel trends and exploring the timing of the effect of policy proposals is as follows:

$$y_{ist} = \beta_{\leq t-5} \mathbb{1}(Proposed_{s, \leq t-5}) + \sum_{j=-4}^{j=+4} \mathbb{1}(Proposed_{s, t+j}) + \alpha_{is} + \alpha_t + \\ + \beta_{\geq t+5} \mathbb{1}(Proposed_{s, \geq t+5}) + Age_{ist} + \delta Z_{st-1} + \varepsilon_{it} \quad (2)$$

where $\mathbb{1}(Proposed_{s, t+j})$ is an indicator for being j periods away from the time of treatment at t . I group years that are five or more years away from the policy proposal. I also normalize the year before the policy proposal to 0.

B. Do policy proposals affect startup fundraising outcomes?

In this section I study the effect of AI regulation proposals on startup fundraising outcomes. The effect of regulation proposals can be understood similarly to announcement effects in financial markets, where they create initial uncertainty and concern among businesses and investors about the potential implications of the upcoming regulations (see for e.g. Bomfim 2003; Jbir et al. 2023).

Insert Table III about here

Columns (1) and (3) in Table III display the difference-in-difference estimates for the effect of AI regulation proposals on startup fundraising outcomes, from estimating equation (1) with a TWFE model as outlined in Section II.A. Column (1) focuses on the effect of proposals on the probability of fundraising, whereas column (3) considers the logged amount raised conditional on fundraising success. AI regulation proposals have a negative and significant effect (at 1%) on the probability of fundraising success amounting to a 2.3% annual decrease. Given an unconditional funding rate of 15.6% in the sample, this corresponds to roughly a 15% decline in the likelihood of raising capital in a given year, indicating a meaningful tightening in access to finance. The same does not hold true for the amounts raised, with the coefficient of ‘Proposed’ being positive and significant at 10%, though the significance does not survive in the robustness checks. Figure II displays the 95% confidence intervals for the dynamic

differences-in-differences estimates from equation (2), with Panel A depicting effects for the probability of fundraising and Panel B depicting effects for the logged amount raised, conditional on fundraising success. Consistent with the parallel trends assumption, the panels reveal no evidence of pre-trends. Panel A highlights that the negative effect on fundraising probability becomes notable two years after the proposal, with the effect increasing over time.

Insert Figure II about here

While the estimates in Table III and Figure II reflect the average effect of AI regulation proposals, the impact varies across startups depending on their activities. For example, startups that rely heavily on personal data transactions are more affected by privacy regulations, while those less involved in data sales face a lower impact (e.g., Canayaz et al., 2022).

To measure this variation, I compare each startup's description from Crunchbase with relevant AI regulation summaries from BCLP (2024a), using OpenAI's text-embedding-ada-002 model to generate embeddings for both the startup descriptions and regulatory summaries. I then calculate a cosine similarity score between the two, representing each startup's exposure to the regulations.⁹ Appendix I illustrates the effectiveness of the exposure measure by presenting examples of the two highest and lowest ranked regulations based on similarity scores. Reviewing

⁹ For an example of using cosine similarity to compute a measure of company exposure to regulation, please consult Kalmenovitz (2023).

the descriptions of both the companies and the regulations suggests that the similarity scores capture the relevance of the regulations in a reasonable way, generally identifying which regulations are more or less applicable to the companies.

Given that states introduce multiple regulations over time, I construct a cumulative "Exposure" measure for each startup, averaging its exposure to all regulations proposed in its state up to a given year. For state-year observations without any proposed regulations, the "Exposure" value is set to zero. The "Exposure" measure ranges from 0.6674 to 0.8464, and an average exposure score of 0.7347, conditional on positive exposure to regulation. Columns (2) and (4) of Table III present the results using this "Exposure" variable instead of the "Proposed" dummy. The findings indicate that a higher degree of exposure to AI regulation reduces the probability of fundraising but does not affect the amount raised, conditional on successful fundraising.

C. Robustness Checks

In this section, I assess the robustness of the main effect of AI regulation on fundraising outcomes. First, I apply two alternative methods to estimate equation (1), addressing the potential bias in TWFE that can occur due to staggered treatment and treatment effect heterogeneity. Columns (1) and (2) of Appendix Table A.I present the estimates using the Sun and Abraham (2021) method, while columns (3) and (4) show results from Callaway and Sant'Anna (2021). In both cases, AI regulation proposals continue to have a negative and statistically significant effect at the 1% level on the probability of fundraising, with the estimated

coefficient being nearly double that of TWFE. This suggests that the TWFE estimates may represent a lower bound of the true effect of AI regulation. However, there is no statistically significant effect of AI regulation proposals on the amount raised, conditional on fundraising success.

I next investigate whether the concentrated introduction of policies in 2023 affects my estimation results. Appendix Table A.II presents TWFE estimates after excluding these proposals by restricting the sample to the end of 2022 in Panel A and to the end of 2021 in Panel B. The results for the probability of fundraising in columns (1) and (3) remain very similar. Additionally, as seen in columns (2) and (4), there continues to be no significant effect of AI regulation proposals on the amount raised, conditional on fundraising success.

Appendix Table A.III conducts a placebo test on the timing of the introduction of AI regulation. Panel A shifts the timing of the first AI proposal back by 2 years, whereas Panel B and C shift it by 4 and 6 years, respectively. Encouragingly, there is no significant effect of AI regulation on the probability of fundraising.

Appendix Table A.IV presents a leave-one-state-out placebo test that excludes each state in turn to assess whether the estimated effect is driven by any single state. The estimates remain statistically significant at the 10% level across all exclusions. When California is omitted, the coefficient declines in absolute size relative to the baseline, consistent with the high concentration of AI startups in that state.

Finally, Appendix II examines the effect of policy enactments among states that have proposed AI regulation. Appendix II.A details the methodology, introducing an instrumental-variables (IV) specification that instruments enactment with the partisan composition of the state senate. The results are reported in Table A.V and discussed in Appendix II.B. According to the 2SLS estimates, enactment reduces the annual probability of fundraising by about 3 % points, a larger effect than the roughly 2 % point decline associated with regulation proposals.

D. Exit

AI regulation may also affect startups' exit strategies. As discussed in Section III.A, the regulation lowers the probability of fundraising, potentially making it more difficult for startups to go public (see for e.g. Eldar and Grennan, 2024). Additionally, the high costs of regulatory compliance may make startups more likely to consider being acquired by larger firms that are better equipped to handle these challenges (Ince, 2024).

Insert Table IV about here

Table IV presents TWFE estimates for the probability of an IPO in columns (1) and (2), and the probability of an acquisition in columns (3) and (4). Once a startup reaches one of these exit outcomes (IPO or acquisition), it is dropped from the panel. In this analysis, I focus on treated

startups that had an IPO or were acquired after the AI regulation proposal. The results suggest that while AI regulation has no significant impact on the likelihood of an IPO, it does increase the probability of a startup being acquired by 1.5% points annually. This effect is almost double the unconditional average acquisition rate of 0.7% in the panel.

III. Regulatory Heterogeneity: Procedural vs Constraining Rules

Section II examines the effects of AI regulation on the aggregate. Yet proposed laws differ across states. Section III then introduces a two-part regulatory taxonomy, formalizes the construct and the coding rules (Section III.A), and applies it to test heterogeneous regulatory effects on AI startups (Section III.B).

A. A taxonomy of AI Regulation: Motivation and Definitions

The development of AI systems poses great societal risks and uncertainty (Chao et al., 2023; Branch et al. 2022; Zhu et al., 2024). These range from known risks such as privacy and fairness (Comunale and Manera, 2024), to unknown risks yet to unfold as the technology develops further (UK Government, 2023; National Institute of Standards and Technology, 2023). In face of these risks, there is international consensus that regulation is necessary to safeguard human rights and other public interests (United Nations, 2024; UNESCO, 2021; UK Government, 2023). Yet, just as there are unforeseen risks from AI technologies, several interest groups worry about unforeseen risks of regulation itself. The most recent example includes the contentions towards the EU AI act, where several interest groups and global personalities have

called for its pause due to potential threats to AI efforts and innovation (Moens & Bradshaw, 2025; Mukherjee, 2025).¹⁰ While the discussion often takes opposing views driven by opposing fears, the question that remains unanswered is not whether to regulate AI, but how to do so.

This is not the first time societies have confronted transformative technologies with serious downside risks. Nuclear power raised persistent safety hazards; the Industrial Revolution produced profound labor-market dislocations; and recombinant-DNA research in the 1970s sparked biosafety concerns that led to a temporary moratorium and the adoption of formal biosafety guidelines (Perrow, 1984; Acemoglu & Restrepo, 2019; Berg et al., 1975). The AI context, however, is distinct: it functions as a general-purpose technology diffusing across sectors; it is highly recombinable, with upstream foundation models cascading into diverse downstream applications; and it is opaque and scale-sensitive, making capabilities and externalities hard to predict *ex ante* (Bresnahan & Trajtenberg, 1995; Wiggins & Tejani, 2022; Burrell, 2016; National Institute of Standards and Technology, 2023).

These features of AI systems create two governance deficits. The first is an **information deficit**: because modern AI systems are complex, adaptive, and often opaque, external decision-makers lack credible, verifiable evidence about how they are built, tested, updated, and how they behave in context; as a result, claims about safety, accuracy, or fairness cannot be reliably assessed (Burrell, 2016; Doshi-Velez & Kim, 2017; National Institute of Standards and

¹⁰ Contemporary reporting documents these innovation-focused objections. On 3 July 2025, the Financial Times reported an open letter by 44 European CEOs (including Airbus, ASML, SAP, and Mistral) urging the Commission to pause the AI Act's application, warning that missing guidance and legal uncertainty could harm competitiveness and innovation (Moens & Bradshaw, 2025). Reuters coverage the same day recorded parallel calls—from CCIA Europe and major U.S. and EU firms including Alphabet, Meta, Mistral and ASML (Mukherjee, 2025).

Technology, 2023; European Commission, 2025; Raji et al., 2020). The second is a **boundary deficit**: because AI capabilities travel across sectors, can be repurposed to harmful ends with minimal tweaks, and intensify with scale, case-by-case discretion is inadequate; clear **ex ante** lines and predictable triggers are needed to guide conduct and enable timely intervention (European Commission, 2025; OECD, 2019/2024; European Parliament and Council of the European Union, 2024). Classical regulatory theory points to two families of tools that map onto these deficits. One family addresses the information deficit by structuring how organizations generate and demonstrate risk knowledge—planning, documentation, audits, and learning (Coglianese & Lazer, 2003; Black, 2001). The other addresses the boundary deficit by drawing hard limits on conduct—clear scope rules and prohibitions, backed by credible escalation mechanisms (Kaplow, 1992; Ayres & Braithwaite, 1992; Coglianese & Lazer, 2003). In this paper I call the first family of regulations **procedural** and the second **constraining**.

Guided by this distinction, I organize the discussion around the two families of rules. I begin with procedural regulations because they directly addresses the information deficit and, in current state proposals, is the modal approach. Conceptually, procedural rules do not outlaw particular uses; they make systems legible over time by requiring actors to turn private development work into stable, reviewable claims that can be checked and acted upon. These claims circulate within firms and to regulators, supporting due diligence and enabling targeted intervention when signals warrant. I then turn to constraining regulation, which operationalizes the boundary deficit by specifying who and what falls within scope, when heightened duties

attach, and which practices are ruled out. Reading the statutes through this lens clarifies mechanisms and underpins the heterogeneity tests that follow.

This procedural – constraining taxonomy is also consistent with how recent work characterizes the channels through which AI regulation affects firms. Kretschmer et al. (2023) describe ex-ante restrictions that set hard limits on which AI systems may be developed or deployed and expose firms to shutdown or enforcement risk if they operate outside those boundaries. These rules map onto what I term constraining regulation: they address the boundary deficit by specifying what is and is not permissible. By contrast, Tartaro et al. (2023) emphasize the role of transparency and standardized obligations in enabling the safe uptake of AI. This perspective aligns with what I term procedural regulation, which targets the information deficit by requiring documentation, auditing, and disclosure rather than prohibiting particular lines of development. In this sense, the taxonomy is not only descriptive of statutory language; it captures two distinct regulatory logics that theory predicts should generate different firm-level responses, particularly for startups.

I treat procedural regulation as a *governance cycle that turns private engineering work into public accountability*. Panel A of Figure III and Table A.IV provide illustrations, definitions, and justifications for each procedural clause. During design and training, contemporaneous documentation creates a verifiable record of data, training runs, evaluations, and incidents (National Institute of Standards and Technology, 2023). Before deployment, an impact assessment translates that record into context-specific judgments and a monitoring plan Portions

of the record then move outward to the regulators through targeted regulatory filings (H. 1873, Massachusetts General Court, 2023). Independent checks in the form of external audits or certification compare paper to practice (H. 6286, Illinois General Assembly, 2023), and public registries link systems and uses to responsible parties (S.31, Massachusetts General Court, 2023). The loop does not freeze design; it reduces information asymmetry, builds buyer trust, and gives regulators sightlines to intervene when warranted.

I treat constraining rules as the boundary architecture that determines *who* is covered, *when* obligations bite, and *what* is off-limits—so that oversight does not depend on discretion alone. Panel B of Figure III and Table A.VI display illustrations, definitions, and accompanying justifications for constraining rules. Coverage starts with the legal definition of an “AI system”, the gate that brings actors and tools into scope (H.B. 49, Pennsylvania General Assembly, 2023). Within that scope, risk thresholds - by capability, context, or domain - switch on stricter duties where harms concentrate (H. 7111, Vermont General Assembly, 2023). To address acute danger, statutes provide emergency brakes (temporary injunctions or suspensions) and, where mitigation is not credible, categorical bans (e.g., certain facial or emotion recognition; B. 25-0114, Council of the District of Columbia, 2023; H. 114, Vermont General Assembly, 2023). Because openness changes who can adapt systems, some bills state how open-source models are treated, so obligations attach predictably (A. 8195, New York State Assembly, 2023). Together these constraints deliver broad but risk-calibrated coverage and hard stops for unacceptable uses, complementing the procedural cycle and giving startups predictable guardrails.

I apply this taxonomy to the full text of each proposed AI regulation, coding the presence or absence of each clause based on a close reading. Table A.VI shows the prevalence of each regulation type. Two features stand out. Procedural obligations are the default approach. Across 23 states, 56 laws include at least one procedural duty; disclosures to regulators (50) and impact assessments (47) are the most common, while third-party audits (9) and registration (7) are used more sparingly. This pattern suggests states are prioritizing tools that generate verifiable evidence and scale across use cases before turning to heavier assurance mechanisms. For startups, that choice lowers fixed compliance costs while still making products legible to customers and supervisors.

Constraining provisions are less common and more targeted. Many laws adopt broad AI definitions (30 laws in 13 states), ensuring coverage of both rules-based and learning-based systems. Far fewer create explicit high-risk tiers (3 laws in 2 states), so calibration often happens through the procedural duties above. Where hard lines exist, they address specific harms—prohibited use cases (13 laws in 8 states)—and provide temporary injunction powers (13 laws in 7 states) to halt deployments that pose significant risks. Open-source treatment appears in only two states.

B. Heterogeneous effect of regulation on AI startups

The previous section introduced two distinct types of AI regulation, procedural and constraining. In this section, I examine how each may affect startup outcomes differently.

Procedural regulations mostly operate through additional processes and documentation that targets the information gap surrounding AI technologies. The net impact of these regulations is ambiguous. On the one hand, the additional requirements are likely to introduce new fixed costs, which can disadvantage under-resourced startups. Prior work shows that procedural burdens affect firm outcomes: Klapper et al. (2006) find that bureaucratic entry regulations shape startup size, and Ewens et al. (2024) quantify disclosure and governance costs, documenting bunching around regulatory thresholds. On the other hand, these processes may reduce information asymmetry by standardizing how claims are documented and disclosed, giving stakeholders clearer, testable information about system quality and risk (National Institute of Standards and Technology, 2023). Because AI systems are unusually opaque to outside evaluators (National Institute of Standards and Technology, 2023), narrowing these information gaps can ease due-diligence frictions for investors and buyers and, where uncertainty is the binding constraint, help support access to capital (Diamond & Verrecchia, 1991). From this perspective, procedural rules can also help establish standards that enhance the safety and trustworthiness of AI systems (Tartaro et al., 2023).

Constraining rules operate mainly through creating boundaries around what startups can develop or how they deploy technology. Broad-scope provisions raise the likelihood that a product falls in-scope (and thus the risk of non-compliance), while temporary injunctions create discontinuous shutdown risk. In combination, these features could heighten policy and enforcement uncertainty, lowering the option value of immediate experimentation and

encouraging delay or scale-back of investment—a pattern consistent with real-options theory and evidence that policy uncertainty depresses capital spending and project starts (Bernanke, 1983; Bloom, 2009; Baker, et al., 2016; Julio & Yook, 2012; Gulen & Ion, 2016).

To test these hypotheses, I re-estimate the equation (1) in Section II, interacting the binary policy-proposal shock with indicators for whether a state has a procedural or a constraining regulation by the end of the sample. Table V reports the results. Columns (1) to (3) display the effect of each regulation type on the probability of raising funding, while columns (4) to (6) display the effect on the probability of being acquired. In Appendix Table A.VII I show that results remain robust to specifying a dynamic version of procedural and constraining regulation, where the dummy switches to 1 in the first year in which the state adopts the respective regulation.

Insert Table V about here

The estimates show that constraining rules amplify the adverse impact of AI regulation on startups. Columns (2) and (5) reveal that startups in states which adopt constraining rules have a lower probability of raising funding and a higher probability of being acquired, compared to treated startups in states without constraining regulations. This pattern aligns with real-options predictions and evidence that policy uncertainty depresses investment (Bernanke, 1983; Bloom,

2009; Baker et al., 2016; Julio & Yook, 2012; Gulen & Ion, 2016). By contrast, procedural rules are offsetting, consistent with information-and-disclosure mechanisms easing due-diligence frictions (Diamond & Verrecchia, 1991; Tartaro et al., 2023). Taken together, the results show that it is regulatory design - not the mere presence of regulation - that drives startup finance and exit outcomes.

IV. Regulation, Governance, and the Organization of Technical Work

Having documented the market-level effects of AI regulation on financing and exits, I now examine how startups adjust their **technical strategy** in response. Section IV.A sets out the conceptual framework and hypotheses; Section IV.B details the data, measurement, and empirical approach; and Section IV.C reports the results.

A. Conceptual framework and hypothesis

At the center of the debate surrounding AI regulation lies a fundamental tension: protecting society from AI risks while preserving the trajectory of innovation. This section addresses that tension by examining how startups reallocate scarce technical resources across these competing objectives. I assess these reallocations through the composition of demand for technical expertise within the firm - distinguishing among capabilities devoted to risk mitigation (and hence regulatory compliance), those focused on bringing existing AI technologies to the market, and those directed toward advancing the frontier through novel methods.

Understanding these reallocations requires situating them within established frameworks of organizational learning and innovation under constraint. Insights from the broader innovation literature suggest that shifts in the external environment—such as regulatory changes—can alter the firm’s innovation payoff landscape, prompting reoptimization of internal effort (see Aghion et al., 2023). A standard benchmark in organizational learning models this adjustment on the exploration–exploitation margin - allocating effort between search/experimentation and refining existing technologies (Levinthal & March, 1981; March, 1991). In a complementary strand, models of dynamic experimentation describe an intertemporal allocation between a safe arm with known, near-term payoffs and a risky, information-generating arm (Bolton & Harris, 1999; Keller et al., 2005). These frameworks highlight how tolerance for early failure shapes firms’ willingness to pursue exploratory projects. When the environment penalizes early-stage failure or delays rewards to innovation, the relative payoff to experimentation declines, leading firms to reallocate effort toward safer, nearer-term projects (Manso, 2011; Ederer & Manso, 2013; Tian & Wang, 2014). Regulatory constraints in AI development may operate through similar channels: as they raise the costs or uncertainty associated with high-risk applications (Tartaro et al., 2023), they can indirectly reduce the attractiveness of exploratory initiatives and shift resources toward lower-risk deployments.

Building on these insights, prior work links institutional and policy frictions to systematic shifts in innovative effort. Employment protection and product-liability exposure, for instance, have been shown to affect technology choice and the allocation of R&D across projects

(Samaniego, 2006; Bartelsman et al., 2016; Galasso & Luo, 2022). Relatedly, threshold-like regulatory frictions can act as an implicit tax on growth, lowering effective returns to R&D and biasing the internal project mix (Aghion et al., 2023). In AI-intensive firms, mandated risk-mitigation measures introduce governance and compliance tasks that scale with activity (Tartaro et al., 2023), effectively raising the cost of high-exposure applications and, in line with these frameworks, may redirect technical effort toward nearer-term or lower-risk deployments. This discussion motivates the following hypotheses:

Hypothesis 1 (Governance orientation). AI regulation increases the allocation of technical resources to governance and compliance-related activities.

Hypothesis 2 (Exploratory effort). AI regulation decreases the allocation of technical resources to exploratory, frontier-advancing activities.

Section III documented variation in financing outcomes across regulatory designs, with constraining and procedural rules operating through distinct mechanisms. Building on that evidence, this section further examines how such design differences may influence firms' internal organization of technical work. Procedural rules emphasize standardized documentation and disclosure processes intended to clarify compliance expectations and address informational gaps surrounding AI systems. Constraining rules, in contrast, define the boundaries of permissible development and deployment, introducing enforcement and compliance considerations that interact with firms' innovation choices. Together, these regulatory approaches

alter the context in which startups allocate technical expertise, potentially shaping the balance between governance, market deployment, and exploratory innovation.

Hypothesis 3 (Regulatory design). Variation in regulatory design affects the internal distribution of technical effort, reflecting distinct organizational responses to procedural and constraining rules.

To evaluate these hypotheses, I examine how firms reallocate technical resources in response to regulatory exposure. These shifts are difficult to observe directly: traditional indicators such as financing or exit outcomes provide only coarse signals and cannot disentangle whether regulation alters exploratory effort, compliance capacity, or near-term deployment. To capture these mechanisms more precisely, I use job postings as a revealed-preference measure of labor demand. Postings indicate where the marginal unit of headcount is allocated - across exploration, exploitation, and governance - offering a window into how firms rebalance their technical resources under regulatory pressure.

B. Measuring AI exploitation, AI exploration and AI governance

Regulatory changes can alter how firms allocate technical effort across different modes of activity. To capture these adjustments, I draw on the organizational learning literature, which distinguishes between *exploration* - search, experimentation, and the pursuit of novel methods - and *exploitation* - the refinement, scaling, and commercialization of existing technologies (Levinthal & March, 1981; March, 1991). In the modern AI context, a third domain has become

increasingly central to deployability: *governance*. This category encompasses the evaluative, documentation, and assurance functions that make AI systems auditable, accountable, and contractible (Mitchell et al., 2019; Raji et al., 2020). This tripartite decomposition - exploration, exploitation, and governance - provides a useful lens for analyzing firms' responses to regulation and financing conditions. Each activity differs in its cash-flow timing, option value, and degree of measurability or contractibility, as well as in its exposure to regulatory oversight. These are precisely the dimensions along which financing and policy shocks are expected to exert their strongest effects (Hall & Lerner, 2010; Kerr et al., 2014; Nanda & Rhodes-Kropf, 2017; Holmström & Tirole, 1997).

To observe these reallocations at the point when they are decided, I use job postings as a revealed-preference measure of marginal capability acquisition - that is, where the next unit of headcount is intended to go. Vacancies are widely employed to study firm labor demand and the task composition of work (e.g., Acemoglu et al., 2022a; Babina et al., 2024b; Davis et al., 2013), and posting text provides rich information about evolving skill requirements, technological focus, and organizational design (e.g., Hershbein & Kahn, 2018; Deming & Kahn, 2018). Compared with downstream outcomes such as product launches or funding events, job postings offer a more immediate view of intended inputs, capturing how firms adjust their technical mix in response to external shocks. This feature makes them well suited for testing whether regulation coincides with changes in the balance of exploration, exploitation, and governance.

Prior work has used job posting data to proxy capability investment at broader levels - for example, firm-level AI job measures linked to growth and product innovation (Babina et al., 2024b), R&D posting activity and skill mixes linked to future innovation (Du et al., 2024), and posting-based “green-skill” intensity to capture investment in green human capital (Darendeli & Law, 2022). Unlike these studies, I decompose AI vacancies into exploration, exploitation, and AI-governance roles and track their within-AI composition at the firm level. This approach turns reallocation into an observable choice: how a continuing startup allocates the next unit of technical headcount across risky experimentation, productization, and governance.

For job postings, I use data from Lightcast (formerly Burning Glass Technologies). The dataset aggregates postings from company career pages and more than 40,000 online job boards, parses free text into structured fields, and applies a deduplication algorithm to avoid counting the same vacancy multiple times. Each record contains the employer name, posting date, job title, job location and full posting text. These features - broad coverage, standardized fields, and deduplication - are well documented in prior studies using the same source (e.g., Babina et al., 2024b).

I link postings from 2012 to 2023 to the startup-year firm panel using exact name matching. To ensure consistent coverage, I retain only startup-year observations that occur after the firm’s first recorded posting, which avoids measurement error from coding missing early postings as zero activity. Internships and postings without job descriptions are excluded before constructing all measures. Finally, I restrict the treated sample to startups with at least one

posting before regulation, ensuring that the estimated effects reflect changes in labor demand rather than initial entry into the data.

I classify postings into three mutually exclusive primary roles (AI exploitation, AI exploration and Non-AI) based on the dominant task profile, and record an orthogonal AI-Governance indicator that may apply to any posting (including Non-AI roles). AI exploitation covers roles whose core deliverables include the refinement, scaling, and operation of existing ML/AI capabilities using established methods. AI exploration includes roles whose core deliverables involve search and experimentation to create new AI/ML technical knowledge, models, or training/optimization schemes; when a posting contains both exploitation and exploration cues, it is classified as AI exploration.¹¹ Postings that satisfy neither set of primary criteria are coded as Non-AI. Throughout the analysis, I use the terms Applied AI and Exploitation interchangeably, and likewise for Frontier AI and Exploration.

The AI-Governance indicator is assigned when the posting bears explicit responsibility for activities aimed at mitigating AI-related risks - such as documentation, audit, and oversight - and for ensuring compliance with AI- or data-governance laws, regulations, or standards.

Because governance tasks can appear across functions, this indicator may co-exist with any

¹¹When a posting contains both exploitation and exploration tasks, I code it as Exploration. This choice aligns the discrete label with the empirical object of interest, which is whether firms allocate any effort to exploration at the margin; the presence of exploratory responsibilities satisfies that condition. It also corresponds to the theoretical representation of choices in models of dynamic experimentation, where decisions are described as an allocation between a safe option and a risky, information-generating option; any positive allocation to the risky option constitutes exploration, so coding mixed postings as Exploration is the closest discrete analogue (Bolton & Harris, 1999; Keller et al., 2005). Finally, it addresses a measurement concern: job advertisements routinely include operational language even for research-oriented roles, and assigning mixed postings to Exploitation would undercount exploration and attenuate estimated reallocation effects.

primary classification - for example, a legal or compliance role in Non-AI, an operations role that includes model-risk controls in Exploitation, or a research role focused on AI safety measures and related assurance practices in Exploration. Appendix IV details the taxonomy, inclusion/exclusion criteria, and tie-breakers used in coding.

Unlike skill-based or keyword approaches, which identify listed capabilities, my objective is to capture the intended function of each role - how firms plan to deploy technical effort. Doing so requires interpreting the task intent and context conveyed in the full job description rather than relying on individual skill mentions. To that end, I hand-label 3,132 postings using the definitions in Appendix IV and fine-tune two RoBERTa classifiers on this corpus. RoBERTa, a transformer-based encoder that refines BERT's pretraining, delivers strong benchmark performance across text-classification tasks (Liu et al., 2019) and has become standard in recent finance and economics applications (Huang et al., 2023; Pfeifer & Marohl, 2023; Moreno & Ordieres-Meré, 2025; Cao et al., 2023). It is well suited to this setting because transfer learning performs effectively with limited hand labels (Howard & Ruder, 2018; Gururangan et al., 2020), and contextual subword representations capture the domain-specific terminology common in AI postings (Sennrich et al., 2016; Kudo & Richardson, 2018).

To operationalize the taxonomy, one model classifies postings into the three mutually exclusive categories—Exploration (Frontier AI), Exploitation (Applied AI), and Non-AI—while the second predicts the binary AI-Governance indicator. Both models demonstrate high accuracy, with F1 scores of 0.837 and 0.866, respectively, providing confidence in the reliability of the

automated classification. Appendix V provides additional details on model specification and evaluation.

My sample consists of 2,883 startups matched to 23,814 AI-related job postings classified into Applied AI, Governance, and Frontier AI roles. Panel B of Table A.VIII presents summary statistics, documenting three patterns. First, aggregate volumes are highest for Applied AI (13,847 postings), followed by Governance (6,661 postings) and Frontier AI (3,735 postings). Second, at the firm–year level, the incidence of any posting is 0.165 for Applied AI, 0.079 for Governance, and 0.068 for Frontier AI; conditional on posting, activity tends to occur in discrete hiring episodes, with mean counts of about five for Applied and Governance ($p95 \approx 17$) and about three for Frontier ($p95 = 10$). Third, the within-AI composition leans toward Applied work: the mean frontier-to-applied ratio is 0.25 ($p95 = 1.0$). Overall, AI hiring is organized around concentrated bursts with a predominant applied component, while governance and frontier roles vary across firm-years.

C. Does AI regulation reallocate hiring within startups?

Here I test whether AI regulation induces within-firm reallocation of effort using the job taxonomy in Section IV.B. I begin with Hypothesis 1, examining how regulation affects demand for governance. I re-estimate equation (1) from Section II with governance hiring measures as dependent variables. Panel A of Table VI reports the results: column (1) uses an indicator for any governance posting in the firm–year; column (2) uses the natural log of the number of governance postings, conditional on at least one posting; and column (3) uses the ratio of

governance to non-governance AI postings. Figure A.1 displays the 95% confidence intervals for the dynamic differences-in-differences estimates of AI regulation on these measures.

Insert Table VI about here

Panel A of Table VI reveals that across these outcomes, the coefficient on Proposed is close to zero and statistically insignificant. Panel B explains this null result by disaggregating regulation by design. The interaction with procedural rules is negative and significant in all three columns: the probability of any governance posting declines by 1.43 pp (column 4), the log count of governance postings conditional on posting falls by 0.3028 (column 5) - about a 26% decrease - and the ratio of governance to non-governance AI postings declines by 0.2260 (column 6). In contrast, the constraining interaction is positive throughout: the probability of any governance posting rises by 1.68 pp, the conditional log count increases by 0.2173 - roughly 24% - and the governance/non-governance ratio increases by 0.1835. Given the mean incidence of governance postings (0.0786) and the mean governance/non-governance ratio (0.3042), these effects are economically meaningful. Interpreting the ratio, a higher value reflects a compositional shift in hiring toward governance roles: under constraining rules, governance demand grows faster than non-governance AI hiring, whereas procedural rules are associated with the reverse pattern.

The pattern dovetails with Section III: only constraining provisions are detrimental for fundraising, and here those same provisions are associated with greater governance hiring and a higher governance/non-governance ratio - consistent with heightened enforcement and shutdown risk that pulls headcount toward governance. Procedural rules move in the opposite direction, consistent with standardization reducing the need to expand governance capacity. The average estimates in Panel A therefore mask offsetting responses that depend on regulatory design.

Insert Table VII about here

Next, I ask how regulation shifts the trade-off between AI exploitation and exploration by analyzing the effect of regulation on Applied and Frontier AI hiring. Table A.IX displays the effect of overall regulation, while Table VII breaks-down these aggregate effects by regulation type. Column 1 (3) of both tables displays estimates for the effect of regulation on whether a startup posted an Applied (Frontier) AI Job in the particular year, column 2 (4) on the number of Applied (Frontier) AI Jobs posted, conditional on at least one such posting, and column 5 on the ratio of Frontier to Applied AI jobs. Figure A.2 displays the 95% confidence intervals for the dynamic differences-in-differences estimates of AI regulation on these measures.

Mirroring the governance results, Table A.IX shows no average effect of regulation on either Applied or Frontier AI hiring. Table VII shows that the aggregate null result reflects

offsetting effects across regulatory designs. Under constraining rules, firms tilt toward Applied hiring; under procedural rules, they tilt toward Frontier. Conceptually, constraining provisions raise enforcement/shutdown risk and narrow permissible scope, increasing the option value of postponing capability-intensive frontier investments and favoring nearer-term, modular deployment roles. Procedural provisions standardize documentation and evaluation, lowering process and information frictions and making deeper exploratory work easier to support, hence the relative shift toward Frontier. This design split aligns with Section III, where only constraining rules are detrimental for fundraising.

The margin that moves differs across activities: Applied AI demand adjusts on the extensive margin (column 1: whether any Applied posting occurs), while Frontier AI demand adjusts on the intensive margin (column 4: how many Frontier postings, conditional on posting). This pattern matches underlying production features: relative to applied AI, frontier labs rely on highly specialized talent and compute-intensive infrastructure, evidenced by the rapid escalation in training compute and associated hardware/energy requirements for state-of-the-art models (Amodei & Hernandez, 2018; Strubell et al., 2019; Patterson et al., 2021). Because these inputs are strong complements and partly irreversible, firms typically vary scale rather than toggling the activity on/off (Milgrom & Roberts, 1995).

In summary, regulation does not uniformly raise or lower AI hiring; it reallocates technical effort across governance, exploitation (applied), and exploration (frontier). Disaggregating by design, constraining provisions are associated with more governance hiring

and a lower frontier-to-applied ratio, while procedural provisions are associated with fewer governance postings and a comparatively higher frontier ratio - yielding near-zero average effects when pooled. Adjustment differs by activity: applied hiring moves on the extensive margin (incidence of any posting), whereas frontier hiring moves on the intensive margin (counts conditional on posting). These results map the regulation–innovation tension into observed staffing choices across compliance, market deployment, and frontier exploration.

V. AI regulation and the geographical relocation of work

Section IV examined how firms respond to AI regulation through reallocating their technical effort across governance, exploitation and exploration, balancing AI safety and compliance obligations against innovation incentives. A further question is whether firms adjust not only what kinds of work they emphasize, but where that work takes place.

Differences across jurisdictions are a recurring source of concern in debates over AI regulation. When procedural or constraining rules tighten in one setting, firms may not necessarily discontinue the affected activity; instead, they can move that activity - and the associated hiring - to a place with fewer obligations. This possibility has prompted concern about the emergence of regulatory ‘AI safe havens,’ in which more permissive jurisdictions attract a larger share of AI development and, in doing so, take on greater exposure to AI risk (Tartaro et al., 2023, p.4).

Early behaviour by large firms suggests that AI regulation is not spatially neutral. Meta has delayed rolling out some of its most advanced AI models in the EU, and Apple has likewise postponed the release of new AI features for EU users, both citing regulatory uncertainty (Weatherbed, 2024; Campbell, 2024). Beyond these anecdotes, a growing literature shows that changes in financial and regulatory policy can directly reshape the geography of high-growth firms. Chen & Ewens (2025) show that the Volcker Rule's limits on bank investments in venture capital reduced funding in the most exposed U.S. regions and prompted young firms to relocate toward established venture hubs. Weik et al. (2024) find that VC-backed startups frequently move their headquarters across borders - often toward the United States - and that such moves are strongly associated with access to foreign venture capital. Peng & Zheng (2022) examine whether differences in environmental regulation drive pollution-intensive firms to shift activity toward more weakly regulated "pollution haven" jurisdictions.

To study whether AI regulation affects the geography of startup activity, I track where firms create AI jobs each year. The idea is not to measure formal headquarter relocation, but to capture shifts in where firms place economically meaningful functions in response to regulatory pressures. AI job postings are a useful proxy for firm activity because they indicate where specialized talent is being hired, where technical capabilities are being built, and where future AI-related output is likely to be produced. This approach also lets me quantify not only whether activity moves at all (an extensive margin response), but also how much activity moves (an intensive margin response), by comparing changes in the volume of AI hiring across locations.

For each firm-year, I measure potential geographic shifts in AI activity by counting the total number of AI jobs (both Applied and Frontier AI) the firm posts in U.S. states that remain non-regulated as of the end of the sample period (2023).¹² Importantly, this captures where the firm is trying to build AI capacity, not where it is officially headquartered. Appendix A.VIII displays relevant definitions and summary statistics.

In total, 1,779 of the 17,582 AI jobs in the data (about 10%) are posted in these non-regulated states, and this behaviour is observed for 188 firms, or about 7% of the 2,883 companies in the sample. Although this is a relatively small share of firms, it is important to study for two reasons. First, these are not marginal firms. Companies that place AI jobs in non-regulated states are larger AI employers: on average, a company that shifts AI hiring to non-regulated states posts 44 AI jobs over its lifetime, compared with only 8 AI jobs for firms that never hire in non-regulated states. Second, shifts in AI hiring are forward-looking investments in technical capacity: new AI vacancies reveal where a firm intends to expand high-skill capabilities next, since openings for specialised roles typically increase in locations where firms plan to grow (Davis, et al., 2013).

Table VIII estimates equation (1) using three different dependent variables. The first dependent variable is an indicator for whether a firm posts at least one AI job in a non-regulated state in a given year (columns 1 and 4). The second is the natural logarithm of the number of AI vacancies posted in non-regulated states, conditional on the firm posting at least one such

¹²For control firms, this is measured as AI postings in non-regulated states other than the one where the company is headquartered.

vacancy in that year (columns 2 and 5). The third is the share of AI jobs located in non-regulated states, defined as the number of AI jobs posted in non-regulated states divided by the firm’s total AI job postings in that year; I refer to this measure as “AI Job Flight” (columns 3 and 6). Panel A reports the average effect of AI regulation on these outcomes. Panel B decomposes the effect by type of regulatory proposal. Figure A.3 shows the corresponding dynamic estimates for each dependent variable, obtained from equation (2).

Insert Table VIII about here

Panel A shows that AI regulation proposals are associated with startups shifting AI hiring to non-regulated states. In column (1), the coefficient on *Proposed* is positive and statistically significant: after a proposal, the probability that a firm posts at least one AI job in a non-regulated state increases by about 0.8 percentage points. In column (2), the coefficient is also positive and significant, indicating that, conditional on posting at least one such job, the number of AI vacancies posted outside the regulating state rises. Column (3) shows a positive and significant effect on “AI Job Flight,” the share of a firm’s AI hiring that takes place in non-regulated states. Together, these results imply that proposed AI regulation increases both the

likelihood that a firm hires AI workers outside the regulating state and the intensity with which it does so.

Panel B separates regulation into procedural rules and constraining rules. Column (4) shows that constraining rules are associated with a statistically significant increase in the probability that a firm posts at least one AI job in a non-regulated state, whereas procedural rules are not. This suggests that firms are more likely to open an out-of-state AI presence when the regulation raises enforcement/shutdown risk and narrows permissible scope of AI, rather than when it simply raises compliance or documentation requirements. In columns (5) and (6), the coefficients on constraining rules remain positive for the number and share of out-of-state AI jobs, but they are estimated less precisely. Conceptually, this is consistent with the idea that the clearest and most immediate response to a constraining proposal is to establish any AI hiring activity elsewhere (the decision to “go outside at all”), while differences across regulation types in how intensely firms reweight their overall AI hiring toward those alternative states are harder to isolate statistically.

Taken together, these results show that AI regulation does not only change the composition of firms’ technical effort (Section IV); it also changes the geography of that effort. Startups respond to proposed regulation by shifting AI hiring to states that have not proposed AI rules, and this response is strongest when the proposed rules are directly constraining rather than purely procedural.

VI. Conclusion

The increasing regulatory focus on AI has sparked key concerns about its impact on innovation. This paper explores how AI regulation affects innovation through the startup ecosystem, particularly by influencing access to capital, exit strategies and the organization of technical work.

Leveraging cross-state variation in AI regulation across the U.S., I find that AI regulation proposals reduce the likelihood of startups securing fundraising by 2% points annually. However, this does not affect the total amount raised by startups that successfully secure funding. Additionally, AI regulation proposals increase the probability of startups being acquired by 1.5% points, suggesting that acquisitions may become a preferred strategy to navigate regulatory hurdles. These effects are not uniform across regulatory approaches. They are driven by “constraining” rules - those that restrict what kinds of AI systems can be built or deployed and raise shutdown or enforcement risk - and are much weaker under “procedural” rules that focus on documentation, auditing, and disclosure. This heterogeneity highlights that the design of AI regulation matters: rules that function as ex-ante limits appear most costly for startups.

Beyond finance and exit, I show that regulation changes how startups allocate scarce technical capacity. Under constraining rules, startups shift hiring away from frontier AI research and toward deploying and commercializing existing systems. At the same time, they increase hiring in AI governance and compliance roles, consistent with redirecting engineering resources toward managing regulatory exposure. Finally, I document that startups adjust not only what

work they do, but where they do it. Firms exposed to regulation become more likely to place AI jobs in unregulated states. This pattern is strongest for constraining rules and points to emerging regulatory arbitrage: sensitive AI work moves to lighter-touch jurisdictions.

Taken together, these results point to three implications. First, AI regulation can shape who builds AI by affecting which startups get funded and which are instead absorbed through acquisition. Second, it can steer what gets built by shifting the types of AI development that firms pursue. Third, it can reshape where AI work happens by altering the geography of AI activity. Crucially, the direction of each effect depends on the design of the regulation.

References

- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of economic perspectives*, 33(2), 3-30.
- Acemoglu, D. (2021). *Harms of AI* (Working Paper No. 29247). National Bureau of Economic Research.
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022a). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293-S340.
- Acemoglu, D., He, A., & Le Maire, D. (2022b). *Eclipse of rent-sharing: the effects of managers' business education on wages and the labor share in the US and Denmark*.
- Acemoglu, D., & Lensman, T. (2024). Regulating transformative technologies. *American Economic Review: Insights*, 6(3), 359-376.
- Aghion, P., Bergeaud, A., & Van Reenen, J. (2023). The impact of regulation on innovation. *American Economic Review*, 113(11), 2894-2936.
- Alanoca, S., Gur-Arieh, S., Zick, T., & Klyman, K. (2025, June). Comparing Apples to Oranges: A Taxonomy for Navigating the Global Landscape of AI Regulation. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (pp. 914-937).
- Amodei, D., & Hernandez, D. (2018, May 16). AI and compute. *OpenAI*. Available at: <https://openai.com/index/ai-and-compute/> (Accessed 24 October 2025)
- An Act amending the act of April 9, 1929 (P.L. 177, No. 175) ... providing for artificial intelligence registry*. H.B. 49, Pennsylvania General Assembly. 2023–2024 Regular Session. (2023). <https://www.palegis.us/legislation/bills/2023/hb49>
- An Act concerning health*. H.B. 1002, 103rd Illinois General Assembly. 2023-2024 Session. (2023). <https://www.ilga.gov/documents/legislation/103/HB/10300HB1002.htm>
- An Act Drafted with the Help of ChatGPT to Regulate Generative Artificial Intelligence Models Like ChatGPT*, S. 31, 193rd Massachusetts General Court. 2023-2024 Session. (2023). <https://malegislature.gov/Bills/193/SD1827>
- An Act Preventing a Dystopian Work Environment*. H. 1873, 193rd Massachusetts General Court. 2023-2024 Session. (2023). <https://malegislature.gov/Bills/193/H1873>
- An Act Relating to Commercial Law—General Regulatory Provisions—Generative Artificial Intelligence Models*. H. 6286, Rhode Island General Assembly. January Session 2023. (2023). <https://webserver.rilegislature.gov/BillText23/HouseText23/H6286.pdf>
- An act relating to creating oversight and liability standards for developers and deployers of inherently dangerous artificial intelligence systems*. H. 711, Vermont General Assembly. 2023–2024 Regular Session. (2024). https://legislature.vermont.gov/bill/status/2024/H.711?utm_source=chatgpt.com
- An act relating to restricting electronic monitoring of employees and employment-related automated decision systems*. H. 114, Vermont General Assembly, 2023–2024 Regular Session (2023). <https://legislature.vermont.gov/Documents/2024/Docs/BILLS/H-0114/H-0114%20As%20Introduced.pdf>

- An act to amend the general business law, in relation to requiring the collection of oaths of responsible use from users of certain high-impact advanced artificial intelligence systems*. A. 8105, New York State Assembly, 2023–2024 Regular Sessions. (2023). <https://www.nysenate.gov/legislation/bills/2023/A8105>
- AN ACT to amend the state technology law ...in relation to prohibiting the development and operation of certain artificial intelligence systems (Part D)*. A. 8195, New York State Assembly. 2023–2024 Regular Sessions. (2023). <https://assembly.state.ny.us/leg/?sh=printbill&bn=A8195&term=2023&Text=Y>
- Ayres, I., & Braithwaite, J. (1992). *Responsive regulation: Transcending the deregulation debate*. Oxford University Press, USA.
- Berg, P., Baltimore, D., Brenner, S., Roblin, R. O., & Singer, M. F. (1975). Summary statement of the Asilomar conference on recombinant DNA molecules. *Proceedings of the National Academy of Sciences*, 72(6), 1981-1984.
- Babina, T., Bahaj, S. A., Buchak, G., De Marco, F., Foulis, A. K., Gornall, W., Mazzola, F. & Yu, T. (2024a). *Customer data access and fintech entry: Early evidence from open banking* (No. w32089). National Bureau of Economic Research.
- Babina, T., Fedyk, A., He, A., & Hodson, J. (2024b). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151, 103745.
- Ballotpedia. 2024. *Historical partisan composition of state legislatures*. Available at: https://ballotpedia.org/Historical_partisan_composition_of_state_legislatures (Accessed 17 September 2024).
- Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4), 1593-1636.
- Bartelsman, E. J., Gautier, P. A., & De Wind, J. (2016). Employment protection, technology choice, and worker allocation. *International Economic Review*, 57(3), 787-826.
- Bomfim, A. N. (2003). Pre-announcement effects, news effects, and volatility: Monetary policy and the stock market. *Journal of Banking & Finance*, 27(1), 133-151.
- BCLP. (2024a). *US state-by-state AI legislation snapshot*. Available at: <https://www.bea.gov/datahttps://www.bclplaw.com/en-US/events-insights-news/us-state-by-state-artificial-intelligence-legislation-snapshot.html> (Accessed 29 July 2024).
- BCLP. (2024b). *AI regulation tracker: UK and EU take divergent approaches to AI regulation*. Available at: <https://www.bclplaw.com/en-US/events-insights-news/ai-regulation-tracker-uk-and-eu-take-divergent-approaches-to-ai-regulation.html> (Accessed 7 October 2024).
- Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The quarterly journal of economics*, 98(1), 85-106.
- Black, J. (2001). Decentring regulation: Understanding the role of regulation and self-regulation in a ‘post-regulatory’ world. *Current legal problems*, 54(1), 103-146.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623-685.
- Bolton, P., & Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2), 349-374.

- Bradshaw, T. (2024). UK regulator examines Microsoft and Amazon's AI dealmaking. *Financial Times*. Available at: <https://www.ft.com/content/0597a834-c101-4d01-893d-87c2eca122c9> (Accessed 4 October 2024).
- Branch, H. J., Cefalu, J. R., McHugh, J., Hujer, L., Bahl, A., Iglesias, D. D. C., Heichman, R. & Darwishi, R. (2022). *Evaluating the susceptibility of pre-trained language models via handcrafted adversarial examples*. arXiv preprint arXiv:2209.02128.
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'?. *Journal of econometrics*, 65(1), 83-108.
- Brown, M., & O'Brien, M. (2025). Senate strikes AI provision from the Republican tax bill after uproar. *AP*. Available at: <https://apnews.com/article/ai-regulation-state-moratorium-congress-39d1c8a0758ffe0242283bb82f66d51a> (Accessed 12 November 2025)
- Bureau of Economic Analysis. 2024. *BEA Data*. Available at: <https://www.bea.gov/data> (Accessed 18 August 2024).
- Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big data & society*, 3(1). <https://doi.org/10.1177/2053951715622512>.
- Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2), 200-230.
- Campbell, C. (2024). Digital Markets Act casts uncertainty over EU launch of Apple Intelligence. *Lexology*. Available at: [https://www.lexology.com/library/detail.aspx?g=1e81431e-add4-4f39-bbe6-14fb6e85b4c&#:~:text=Following%20Apple's%20announcement%20last%20month,Digital%20Markets%20Act%20\(DMA\)](https://www.lexology.com/library/detail.aspx?g=1e81431e-add4-4f39-bbe6-14fb6e85b4c&#:~:text=Following%20Apple's%20announcement%20last%20month,Digital%20Markets%20Act%20(DMA)) (Accessed 25 August 2024).
- Canayaz, M., Kantorovitch, I., & Mihet, R. (2022). *Consumer privacy and value of consumer data*. Swiss Finance Institute Research Paper, (22-68).
- Cao, S., Jiang, W., Yang, B., & Zhang, A. L. (2023). How to talk when a machine is listening: Corporate disclosure in the age of AI. *The Review of Financial Studies*, 36(9), 3603-3642.
- CCIA. (2024). *CCIA statement opposing California AI bill as it steps closer to passage*. Available at: https://ccianet.org/news/2024/08/ccia-statement-opposing-california-ai-bill-as-it-steps-closer-to-passage/?utm_source=chatgpt.com (Accessed 12 November 2025)
- Chandler Jr, A. D. (1991). The functions of the HQ unit in the multibusiness firm. *Strategic management journal*, 12(S2), 31-50.
- Chao, P., Robey, A., Dobriban, E., Hassani, H., Pappas, G. J., & Wong, E. (2023). *Jailbreaking black box large language models in twenty queries*. arXiv preprint arXiv:2310.08419.
- Charkoudian, S.G., Tene, O. & Gomez, M. 2024. How states are stepping in to regulate AI. *Goodwin*. Available at: <https://www.goodwinlaw.com/en/insights/publications/2024/09/insights-technology-aiml-how-states-are-stepping-in-to-regulate-ai> (Accessed 3 October 2024)

- Chen, J., & Ewens, M. (2025). Venture capital and startup agglomeration. *The Journal of Finance*, 80 (4), 2153-2198.
- Coglianesi, C., & Lazer, D. (2003). Management-based regulation: Prescribing private management to achieve public goals. *Law & Society Review*, 37(4), 691-730.
- Collis, D., Young, D., & Goold, M. (2007). The size, structure, and performance of corporate headquarters. *Strategic Management Journal*, 28(4), 383-405.
- Comunale, M., & Manera, A. (2024). *The economic impacts and the regulation of AI: A review of the academic literature and policy actions*. International Monetary Fund.
- Criddle & Hammond (2024). California governor vetoes bill to regulate artificial intelligence. *Financial Times*. Available at: <https://www.ft.com/content/b3b92693-a960-4b6c-a503-f2792c77b04d> (Accessed 12 November 2025).
- Cuéllar, M. F., Larsen, B., Lee, Y. S., & Webb, M. (2024). Does information about AI regulation change manager evaluation of ethical concerns and intent to adopt AI?. *The Journal of Law, Economics, and Organization*, 40(1), 34-75.
- Cugurullo, F., & Acheampong, R. A. (2023). Fear of AI: an inquiry into the adoption of autonomous cars in spite of fear, and a theoretical framework for the study of artificial intelligence technology acceptance. *AI & SOCIETY*, 1-16.
- Davis, S. J., Faberman, R. J., & Haltiwanger, J. C. (2013). The establishment-level behavior of vacancies and hiring. *The Quarterly Journal of Economics*, 128(2), 581-622.
- Darendeli, A., Law, K. K., & Shen, M. (2022). Green new hiring. *Review of Accounting Studies*, 27(3), 986-1037.
- Deming, D., & Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1), S337-S369.
- Denes, M., Howell, S. T., Mezzanotti, F., Wang, X., & Xu, T. (2023). Investor tax credits and entrepreneurship: Evidence from us states. *The Journal of Finance*, 78(5), 2621-2671.
- Diamond, D. W., & Verrecchia, R. E. (1991). Disclosure, liquidity, and the cost of capital. *The journal of Finance*, 46(4), 1325-1359.
- Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*
- Du, Y., Wang, W., Zhang, B., & Zheng, Y. (2024). Taking the Pulse of Firm Innovation from Online Job Postings. Available at: <https://ssrn.com/abstract=4707171>
- Eldar, O., & Grennan, J. (2024). Common venture capital investors and startup growth. *The Review of Financial Studies*, 37(2), 549-590.
- Ederer, F., & Manso, G. (2013). Is pay for performance detrimental to innovation?. *Management Science*, 59(7), 1496-1513.
- European Commission. (2025). *Commission publishes the guidelines on prohibited artificial intelligence (AI) practices, as defined by the AI Act*. Available at: [Commission publishes the Guidelines on prohibited artificial intelligence \(AI\) practices, as defined by the AI Act. | Shaping Europe's digital future](#).

- European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union (L), 12 July 2024. <https://data.europa.eu/eli/reg/2024/1689/oj>
- Ewens, M., Xiao, K., & Xu, T. (2024). Regulatory costs of being public: Evidence from bunching estimation. *Journal of Financial Economics*, 153, 103775.
- Fazio, C., Guzman, J., & Stern, S. (2020). The impact of state-level research and development tax credits on the quantity and quality of entrepreneurship. *Economic Development Quarterly*, 34(2), 188-208.
- Fields, J., Chovanec, K., & Madiraju, P. (2024). A survey of text classification with transformers: How wide? how large? how long? how accurate? how expensive? how safe?. *IEEE Access*, 12, 6518-6531.
- Foster, J. (2023). How Do State Legislatures Make Bills and Pass Laws? *Legistorm*. Available at: <https://info.legistorm.com/blog/how-do-state-legislatures-make-bills-and-pass-laws> (Accessed 2 October 2024).
- Frey, C. B., & Presidente, G. (2024). Privacy regulation and firm performance: Estimating the GDPR effect globally. *Economic Inquiry*.
- Future of Life Institute. (2023). *Pause giant AI experiments: An open letter*. <https://futureoflife.org/open-letter/pause-giant-ai-experiments/> (Accessed 10 November 2025)
- Galasso, A., & Luo, H. (2022). When does product liability risk chill innovation? Evidence from medical implants. *American Economic Journal: Economic Policy*, 14(2), 366-401.
- Gans, J. S., Hsu, D. H. and Stern, S. (2002). When Does Start-Up Innovation Spur the Gale of Creative Destruction? *The RAND Journal of Economics*, 33(4), 571-586.
- Gans, J. S., & Stern, S. (2003). The product market and the market for “ideas”: commercialization strategies for technology entrepreneurs. *Research policy*, 32(2), 333-350.
- Gofman, M., & Jin, Z. (2023). Artificial intelligence, education, and entrepreneurship. *The Journal of Finance*, 79(1), 631-667.
- Guha, N., Lawrence, C. M., Gailmard, L. A., Rodolfa, K. T., Surani, F., Bommasani, R., Raji, I. D., Cuéllar, M. F., Honigsberg, C., Liang, P., & Ho, D. E. (2023). *The AI regulatory alignment problem*. Stanford HAI. <https://hai.stanford.edu>
- Guerreiro, J., Rebelo, S., & Teles, P. (2023). *Regulating Artificial Intelligence* (No. w31921). National Bureau of Economic Research.
- Gulen, H., & Ion, M. (2016). Policy uncertainty and corporate investment. *The Review of financial studies*, 29(3), 523-564.
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*.
- Hall, B. H., & Lerner, J. (2010). The financing of R&D and innovation. In *Handbook of the Economics of Innovation* (Vol. 1, pp. 609-639). North-Holland.

- Hammond, G. and Criddle, C.(2024a).. California governor vetoes bill to regulate artificial intelligence. *Financial Times*. Available at:<https://www.ft.com/content/b3b92693-a960-4b6c-a503-f2792c77b04d> (Accessed 1 October 2024).
- Hammond, G. and Criddle, C.(2024b). The AI bill driving a wedge through Silicon Valley. *Financial Times*. Available at:<https://www.ft.com/content/c5e7bf0b-cadc-4673-b2d8-daa9d89f2230> (Accessed 1 October 2024).
- Hershbein, B., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review*, 108(7), 1737-1772.
- Holmstrom, B., & Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *Quarterly Journal of economics*, 112(3), 663-691.
- Howard, J., & Ruder, S. (2018). Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.
- Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2), 806-841.
- Ince, B. (2024). How do regulatory costs affect mergers and acquisitions decisions and outcomes?. *Journal of Banking & Finance*, 163, 107156.
- Jbir, H., Oros, C., & Popescu, A. (2024). Macroprudential policy and financial system stability: an aggregate study. *Empirical Economics*, 66(5), 1941-1973.
- Jia, J., Jin, G. Z., & Wagman, L. (2021). The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science*, 40(4), 661-684.
- Julio, B., & Yook, Y. (2012). Political uncertainty and corporate investment cycles. *The journal of finance*, 67(1), 45-83.
- Kalmenovitz, J. (2023). Regulatory intensity and firm-specific exposure. *The review of financial studies*, 36(8), 3311-3347.
- Kaplow, L. (2013). Rules versus standards: An economic analysis. In: *Scientific Models of Legal Reasoning*. Routledge. pp. 11-84.
- Keller, G., Rady, S., & Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1), 39-68.
- Kerr, W. R., Nanda, R., & Rhodes-Kropf, M. (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives*, 28(3), 25-48.
- Kerr, W. R., & Nanda, R. (2015). Financing innovation. *Annual review of financial economics*, 7(1), 445-462.
- Klapper, L., Laeven, L., & Rajan, R. (2006). Entry regulation as a barrier to entrepreneurship. *Journal of Financial Economics*, 82(3), 591-629.
- Knight, W. & Dave, P. (2023). In Sudden Alarm, Tech Doyens Call for a Pause on ChatGPT. *Wired*. Available at: <https://www.wired.com/story/chatgpt-pause-ai-experiments-open-letter/> (Accessed 10 November, 2025).

- Koski, H., & Valmari, N. (2020). *Short-term Impacts of the GDPR on Firm Performance* (No. 77). ETLA Working Papers.
- Kretschmer, M., Kretschmer, T., Peukert, A., & Peukert, C. (2023). *The risks of risk-based AI regulation: taking liability seriously*. arXiv preprint arXiv:2311.14684.
- Kudo, T., & Richardson, J. (2018). SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*.
- Lee, I. H. (2022). Startups, relocation, and firm performance: a transaction cost economics perspective. *Small Business Economics*, 58(1), 205-224.
- Levinthal, D., & March, J. G. (1981). A model of adaptive organizational search. *Journal of economic behavior & organization*, 2(4), 307-333.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Manso, G. (2011). Motivating innovation. *The journal of finance*, 66(5), 1823-1860.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1), 71-87.
- Maryland General Assembly. (2024). *The legislative process: how a bill becomes a law*. Available at: <https://msa.maryland.gov/msa/mdmanual/07leg/html/proc.html> (Accessed 2 October 2024).
- Milgrom, P., & Roberts, J. (1995). Complementarities and fit strategy, structure, and organizational change in manufacturing. *Journal of accounting and economics*, 19(2-3), 179-208.
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, 220-229.
- Moen, B. & Bradshaw, T. (2025). European CEOs urge Brussels to halt landmark AI Act. *Financial Times*. Available at: [European CEOs urge Brussels to halt landmark AI Act](#)
- Moreno, A., & Ordieres-Meré, J. (2025). Predicting stock price trends using language models to extract the sentiment from analyst reports: Evidence from IBEX 35-listed companies. *Economics Letters*, 112404.
- Mukherjee, S. (2025). Explainer: Will the EU delay enforcing its AI Act? *Reuters*. Available at: [Explainer: Will the EU delay enforcing its AI Act? | Reuters](#)
- Nanda, R., & Rhodes-Kropf, M. (2017). Financing risk and innovation. *Management science*, 63(4), 901-918.
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework*. Available at: <https://doi.org/10.6028/NIST.AI.100-1>
- OpenAI. (2024). Opposition letter on California SB 1047.
- Organisation for Economic Co-operation and Development. (2019). *Recommendation of the Council on Artificial Intelligence* [Adopted May 2019; updated May 2024]. <https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449> (see also overview: <https://oecd.ai/en/ai-principles>)

- OECD, (2024). *OECD AI Incidents Monitor (AIM)*. Available at: https://oecd.ai/en/incidents?search_terms=%5B%5D&and_condition=false&from_date=2014-01-01&to_date=2024-10-11&properties_config=%7B%22principles%22:%5B%5D,%22industries%22:%5B%5D,%22harm_types%22:%5B%5D,%22harm_levels%22:%5B%5D,%22harmed_entities%22:%5B%5D%7D&only_threats=false&order_by=date&num_results=20 (Accessed 11 October 2024)
- Office of the Governor of California. (2024, September 29). *SB-1047 veto message*. Available at: [SFresno Biz24092910250](https://www.sos.ca.gov/legislation/bills/sb100-1099/sb1047-veto-message) (Accessed 12 November 2025)
- Open AI. 2022. *New and improved embedding model*. Available at: <https://openai.com/index/new-and-improved-embedding-model/> (Accessed 5 October 2024).
- Open AI. 2025. *Opposition letter on California SB 1047*. Available at: <https://www.documentcloud.org/documents/25056617-ca-sb-1047-openai-opposition-letter/> (Accessed 12 November 2025)
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L. M., Rothchild, D., ... & Dean, J. (2021). Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.
- Peng, N., & Zhang, X. (2022). The impact of environmental regulations on the location choice of newly built polluting firms: based on the perspective of new economic geography. *Environmental Science and Pollution Research*, 29(39), 59802-59815.
- Perrigo, B. (2023). Elon Musk Signs Open Letter Urging AI Labs to Pump the Brakes. *Time*. Available at: <https://time.com/6266679/musk-ai-open-letter/> (Accessed 10 November 2025).
- Perrigo, B., & Chow, A. R. (2025, July 1). Senators reject 10-year ban on state-level AI regulation, in blow to Big Tech. *TIME*. Available at: [Senators Reject 10-Year Ban on State-Level AI Regulation | TIME](https://www.time.com/2025/07/01/senators-reject-10-year-ban-on-state-level-ai-regulation/) (Accessed 12 November 2025)
- Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. Basic Books.
- Pfeifer, M., & Marohl, V. P. (2023). CentralBankRoBERTa: A fine-tuned large language model for central bank communications. *The Journal of Finance and Data Science*, 9, 100114.
- Phillips, G. M., & Zhdanov, A. (2013). R&D and the incentives from merger and acquisition activity. *The Review of Financial Studies*, 26(1), 34-78.
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *In Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 33-44.
- Safe and Secure Innovation for Frontier Artificial Intelligence Systems Act*. SB 1047. California Legislature. 2023–2024 Regular Session. (2023). https://leginfo.ca.gov/faces/billHistoryClient.xhtml?bill_id=202320240SB1047
- Samaniego, R. M. (2006). Employment protection and high-tech aversion. *Review of Economic Dynamics*, 9(2), 224-241.


- Samuel, J., Khanna, T., & Sundar, S. (2024). *Fear of Artificial Intelligence? NLP, ML and LLMs based discovery of AI-phobia and fear sentiment propagation by AI news. NLP, ML and LLMs Based Discovery of AI-Phobia and Fear Sentiment Propagation by AI News* (March 09, 2024).
- Schutjens, V., & Stam, E. (2003). The evolution and nature of young firm networks. *Small Business Economics*, 21, 115–134.
- Sennrich, R., Haddow, B., & Birch, A. (2015). Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Stop Discrimination by Algorithms Act of 2023. B. 25-0114, Council of the District of Columbia. 25th Council. (2023). <https://lims.dccouncil.gov/Legislation/B25-0114>
- Strubell, E., Ganesh, A., & McCallum, A. (2020, April). Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 34, No. 09, pp. 13693-13696).
- Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2), 175-199.
- Tartaro, A., Smith, A. L., & Shaw, P. (2023). *Assessing the impact of regulations and standards on innovation in the field of AI*. arXiv preprint arXiv:2302.04110.
- Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research policy*, 15(6), 285-305.
- The Guardian. (2024). Meta pulls plug on release of advanced AI model in EU. Available at: <https://www.theguardian.com/technology/article/2024/jul/18/meta-release-advanced-ai-multimodal-llama-model-eu-facebook-owner> (Accessed 25 August 2024).
- Tian, X., & Wang, T. Y. (2014). Tolerance for failure and corporate innovation. *The Review of Financial Studies*, 27(1), 211-255.
- UK Government. (2023). *The Bletchley Declaration by countries attending the AI Safety Summit (1–2 November 2023)*. Available at: <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration>
- UNESCO. (2021). *Ethics of Artificial Intelligence*. Available at: <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>
- United Nations General Assembly. (2024). *Seizing the opportunities of safe, secure and trustworthy artificial intelligence systems for sustainable development*. Available at: <https://docs.un.org/en/A/RES/78/265>
- Weatherbed, J. (2024). Meta won't release its multimodal Llama AI model in the EU. *The Verge*. Available at: <https://www.theverge.com/2024/7/18/24201041/meta-multimodal-llama-ai-model-launch-eu-regulations> (Accessed 25 August 2024).
- Weik, S., Achleitner, A. K., & Braun, R. (2024). Venture capital and the international relocation of startups. *Research Policy*, 53(7), 105031.

- Wheeler, T. (2023). The three challenges of AI regulation. *Brookings Institute*. Available at: <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/> (Accessed 9 October, 2024)
- Wei, K., Ezell, C., Gabrieli, N., & Deshpande, C. (2024, August 20). *How do AI startups "fine-tune" policy? Examining regulatory capture in AI governance*. 2024 AAAI/ACM Conference on AI, Ethics, and Society. SSRN. <https://ssrn.com/abstract=4931927>
- Werthman, C. (2018). Democrats are Becoming the Party of States' Rights. *Brown Political Review*.
- White & Case. (2024). *AI Watch: Global regulatory tracker - United States*. Available at: <https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-united-states#:~:text=Currently%2C%20there%20is%20no%20comprehensive,prohibit%20or%20restrict%20their%20use> (Accessed 2 October 2024)
- Wiggins, W. F., & Tejani, A. S. (2022). On the opportunities and risks of foundation models for natural language processing in radiology. *Radiology: Artificial Intelligence*, 4(4), e220119.
- Wu, Y., & Wan, J. (2025). A survey of text classification based on pre-trained language model. *Neurocomputing*, 616, 128921.
- Zhu, B., Mu, N., Jiao, J., & Wagner, D. (2024). *Generative AI security: challenges and countermeasures*. arXiv preprint arXiv:2402.12617.

Figure 1: Distribution and timing of AI regulation

Panel A of this figure illustrates the distribution of proposed AI regulations across U.S. states as of December 31, 2023, while Panel B shows the timeline of both proposed and enacted AI regulations.

Panel A: Distribution of AI policy proposals across states



Panel B: The timing of AI regulation proposals and enactments





Figure II: Dynamic Effect of AI regulation proposals on firm financing outcomes

This figure illustrates the dynamic effects of AI regulation proposals on startups' fundraising performance, as estimated by equation (2) in Section II.A. The dots represent point estimates of the dynamic difference-in-differences coefficients, while the bars show the 95% confidence intervals. The year prior to the regulation proposal is normalized to zero. In Panel A, the dependent variable is a binary indicator for whether a startup raised funding in a given year. In Panel B, the dependent variable is the natural logarithm of the amount raised, conditional on successful fundraising. The model controls for the lagged log level and growth rate of real GDP, lagged log disposable income per capita, startup fixed effects, year fixed effects, and age fixed effects. Standard errors are clustered at the state level.

Panel A: Raised Eq



Panel B: Amount Raised (log)




Figure III

These figures illustrate the AI regulation taxonomy, as described, defined and explained in Section III and Table A.VI. Panel A depicts procedural regulation clauses, while Panel B shows constraining regulation clauses.

Panel A: Procedural Regulations



Panel B: Constraining Regulations



Table I: Variable Definitions

Variable Definitions

Dependent Variables

Raised Eq	A dummy variable equal to 1 if the company raised investment each calendar year since incorporation. Source: Crunchbase.
Annual Inv	The natural logarithm of the annual amount raised by a particular company conditional on fundraising success, in million USD Source: Crunchbase.
IPO	A dummy variable equal to 1 if the company had an IPO in each calendar year. Once a company has an IPO, it is dropped from the sample. Source: Crunchbase.
Acquired	A dummy variable equal to 1 if the company was acquired in each calendar year. Once a company is acquired, it is dropped from the sample Source: Crunchbase.

Independent Variables

Proposed	A dummy variable equal to 1 if the state in which the company is located proposed its first AI regulation in a particular year. Dummy remains 1 for all years subsequent to initial proposal. Source: BCLP.
Exposure	The average cosine similarity between a company's description and the summaries of AI regulations proposed in the company's state up to a given year. This similarity is calculated using embeddings generated by OpenAI's text-embedding-ada-002 model. Source: BCLP&Crunchbase.
Lag Real GDP	The natural logarithm of real GDP (2017 dollars) in the state the company is located in, in the previous year. Source: Bureau of Economic Analysis, U.S. Department of Commerce (BEA).
Lag Real GDP Growth Rate	Growth rate of real GDP (2017 dollars) in the state the company is located in, in the previous year. Source: Bureau of Economic Analysis, U.S. Department of Commerce (BEA).
Lag Disposable Income p.c.	The natural logarithm of disposable personal income per capita (current dollars) in the state the company is located in, in the previous year. Source: Bureau of Economic Analysis, U.S. Department of Commerce (BEA).

Fixed Effects

Company	A set of dummy variables for each company. Source: Crunchbase.
Year	A set of dummy variables for each calendar year. Source: Crunchbase.
Age	A set of dummy variables for the age of the company in the particular calendar year. Source: Crunchbase.

Table II: Summary Statistics

This table presents summary statistics at both the startup-year and state-year levels, covering 13,463 startups across the U.S. Specifically, there are 10,957 treated startups in 28 states with AI regulation proposals and 2,506 control startups in states without such proposals. For Table II, the following definitions apply: Annual Inv (USD mill) refers to the annual amount raised by a startup (in million USD), conditional on successful fundraising. Other startup-year variables are defined as in Table I. The number of observations for IPO and Acquired decreases for two reasons: once startups achieve either outcome, they are removed from the sample, and startups in treated states that achieved either outcome prior to the treatment are excluded. State-year variables are defined contemporaneously; for instance, Real GDP (USD billion) reflects each state's real GDP (in billion USD) for the given year.

	Obs.	Mean	St. Dev.	p5	p50	p95
<i>Company-year level</i>						
Raised Eq	99,874	0.1555	0.3624	0.0000	0.0000	1.0000
Annual Inv (USD million)	11,729	16.3921	143.4527	0.0507	2.1000	55.0000
Exposure	99,875	0.2608	0.3517	0.0000	0.0000	0.7544
IPO	99,524	0.0004	0.0200	0.0000	0.0000	0.0000
Acquired	94,589	0.0069	0.0828	0.0000	0.0000	0.0000
<i>State-year level</i>						
Proposed	612	0.0801	0.2716	0.0000	0.0000	1.0000
Enacted	612	0.0539	0.2260	0.0000	0.0000	1.0000
Real GDP (USD billion)	612	386.3473	491.7804	48.6101	224.9411	1,559.6360
Disposable Income p.c.	612	46,864.9400	9,364.5880	33,558	45,797	64,656
Real GDP Growth Rate	612	1.8545	2.5506	-2.5743	1.9139	5.6614

Table III: AI regulation proposals and fundraising outcomes

This table presents the difference-in-differences estimates of the effect of AI regulation proposals on startup fundraising outcomes. Columns (1) and (3) show estimates from the Two-Way Fixed Effects (TWFE) model, based on Equation (1) in Section II.A. The variable "Proposed" is a dummy for state-year observations starting from the year when the first AI regulation is proposed. Columns (2) and (4) provide a variation of the TWFE model, replacing the proposal dummy with a continuous measure of regulation proposals. The variable "Exposure" measures the average cosine similarity between a startup's description and the summaries of AI regulations proposed in the startup's state up to a given year. This similarity is calculated using embeddings generated by OpenAI's text-embedding-ada-002 model. In columns (1) and (2), the dependent variable is a binary indicator of whether a startup raised funding in a given year. In columns (3) and (4), the dependent variable is the natural logarithm of the amount raised, conditional on successful fundraising. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Because of the inclusion of multi-dimensional fixed effects, 145 singleton observations are dropped in columns (1) and (2), and 2,362 are dropped in columns (3) and (4). Standard errors are clustered at the state level, and significance levels are indicated as follows: ***, **, and * for 99%, 95%, and 90% confidence, respectively.

	DV: Raised Eq		DV: Amount Raised (log)	
	(1)	(2)	(3)	(4)
Proposed	-0.0231*** (0.0056)		0.0775* (0.0436)	
Exposure		-0.0302*** (0.0077)		0.0919 (0.0571)
Time-varying state controls	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	99,729	99,729	9,367	9,367
Adjusted R ²	0.2479	0.2479	0.7089	0.7089

Table IV: AI regulation proposals and startup exit

This table presents the difference-in-differences estimates of the effect of AI regulation proposals on startup exit outcomes. Columns (1) and (3) show estimates from the Two-Way Fixed Effects (TWFE) model, based on Equation (1) in Section II.A. The variable "Proposed" is a dummy for state-year observations starting from the year when the first AI regulation is proposed. Columns (2) and (4) provide a variation of the TWFE model, replacing the proposal dummy with a continuous measure of regulation proposals. The variable "Exposure" measures the average cosine similarity between a startup's description and the summaries of AI regulations proposed in the startup's state up to a given year. This similarity is calculated using embeddings generated by OpenAI's text-embedding-ada-002 model. In columns (1) and (2), the dependent variable is a binary indicator of whether a startup had an IPO. In columns (3) and (4), the dependent variable is a binary indicator for whether the startup was acquired. Once a startup has the respective exit outcome, it is dropped from the panel. Normally, I only retain treated startups which had an IPO/were acquired after the regulation proposal. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Because of the inclusion of multi-dimensional fixed effects, 145 singleton observations are dropped in columns (1) and (2) and 150 observations are dropped in columns (3) and (4). Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

	DV: IPO		DV: Acquired	
	(1)	(2)	(3)	(4)
Proposed	0.0004 (0.0002)		0.0145*** (0.0014)	
Exposure		0.0005 (0.0003)		0.0195*** (0.0018)
Time-varying state controls	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	99,379	99,380	94,439	94,439
Adjusted R ²	0.0333	0.7091	0.0528	0.0528

Table V: Heterogeneous effects of procedural versus constraining AI regulation on startup fundraising and exit outcomes.

This table reports difference-in-differences estimates of the heterogeneous effects of AI regulation proposals on startup fundraising and exit. “Proposed” equals 1 for state–year observations beginning in the first year a state proposes an AI regulation. “Procedural Regulations” and “Constraining Regulations” indicate whether, by the end of 2023, a state has proposed a procedural or a constraining regulation, respectively. The dependent variable in columns (1) - (3) is a binary indicator of whether a startup has raised funding. The dependent variable in columns (4) - (6) is a binary indicator for whether the startup was acquired. Startups exit the panel upon acquisition; for columns (4)–(6), the sample retains only treated startups whose acquisition occurs after the proposal date. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

	DV: Raised Eq			DV: Acquired		
	(1)	(2)	(3)	(4)	(5)	(6)
Proposed	-0.0393*** (0.0096)	0.0121 (0.0104)		0.0187*** (0.0025)	0.0075** (0.0029)	
Proposed#Procedural Regulations	0.0275* (0.0162)		0.0216* (0.0109)	-0.0071** (0.0031)		-0.0016 (0.0028)
Proposed#Constraining Regulations		-0.0413*** (0.0112)	-0.0411*** (0.0064)		0.0083** (0.0034)	0.0157*** (0.0026)
Time-varying state controls	Yes	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	99,729	99,729	99,729	94,439	94,439	94,439
Adjusted R ²	0.2481	0.2482	0.2483	0.0530	0.0530	0.0529

Table VI: AI regulation proposals and governance hiring

This table reports difference-in-differences estimates of the effect of AI regulation proposals on startup governance. Panel A shows overall effects; Panel B shows heterogeneity by regulation type. “Proposed” equals 1 from the first year a state proposes an AI regulation. “Procedural” and “Constraining” indicate whether, by end-2023, a state has proposed a procedural or a constraining measure. Dependent variables are: any AI-governance posting (columns 1, 4), the number of such postings conditional on at least one (columns 2, 5), and the ratio of AI-governance postings to the number of non-governance AI Jobs (Applied and Frontier AI Jobs) (columns 3, 6). Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

Panel A: Effect of Overall Regulation on Governance Hiring

	DV: Posted Governance Job (1)	DV: Nr. Governance Jobs (logged) (2)	DV: Ratio Governance/Non- Governance AI Jobs (3)
Proposed	0.0083 (0.0074)	-0.0373 (0.0835)	0.0235 (0.0804)
Time-varying state controls	Yes	Yes	Yes
Company FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Observations	16,669	1,061	2,596
Adjusted R ²	0.3608	0.4678	0.3051

Panel B: Heterogeneous Effect of Regulation on Governance Hiring

	DV: Posted Governance Job (4)	DV: Nr. Governance Jobs (logged) (5)	DV: Ratio Governance/Non- Governance AI Jobs (6)
Proposed#Procedural Regulations	-0.0143** (0.0069)	-0.3028*** (0.1082)	-0.2260*** (0.0747)
Proposed#Constraining Regulations	0.0168* (0.0093)	0.2173** (0.0927)	0.1835*** (0.0659)
Time-varying state controls	Yes	Yes	Yes
Company FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Observations	16,669	1,061	2,596
Adjusted R ²	0.3608	0.4711	0.3600

Table VII: AI regulation proposals and strategic pivoting

This table reports difference-in-differences estimates of the effect of AI regulation proposals on startup Applied & Frontier AI hiring. “Proposed” equals 1 from the first year a state proposes an AI regulation. “Procedural” and “Constraining” indicate whether, by end-2023, a state has proposed a procedural or a constraining measure, respectively. Dependent variables are: any Applied (Frontier) AI posting in column 1 (3), the number of Applied (Frontier) AI postings conditional on at least one such posting in column 2 (4) and the ratio between Frontier and Applied AI postings in column 5. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

	DV: Posted Applied AI Job	DV: Cond. Nr. Applied AI Jobs (logged)	DV: Posted Frontier AI Job	DV: Cond. Nr Frontier AI Jobs (logged)	DV: Ratio Frontier/Applied AI Job
	(1)	(2)	(3)	(4)	(5)
Proposed#Procedural Regulations	-0.0211* (0.0110)	0.0030 (0.0824)	-0.0033 (0.0060)	0.1526* (0.0755)	0.4850** (0.2328)
Proposed#Constraining Regulations	0.0233** (0.0105)	-0.0553 (0.0789)	0.0073 (0.0077)	-0.2181** (0.0873)	-0.3463* (0.1761)
Time-varying state controls	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes
Observations	16,669	2,270	16,669	831	2,270
Adjusted R ²	0.3362	0.4253	0.3153	0.4249	0.1905

Table VIII: AI regulation proposals and relocation of startup activity

This table reports difference-in-differences estimates of AI regulation proposals on relocation of AI hiring. Panel A shows overall effects; Panel B shows heterogeneity by regulation type. “Proposed” equals 1 from the first year a state proposes an AI regulation. “Procedural” and “Constraining” indicate whether, by end-2023, a state has proposed a procedural or a constraining measure. Dependent variables are: any AI posting in non-regulated states, where is firm is not headquartered (columns 1, 4), the number of such postings in non-regulated states conditional on at least one posting (columns 2, 5), and the ratio of AI postings in non-regulated states to total AI vacancies (columns 3, 6). Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

Panel A: Effect of AI regulation on Out-of-state AI Job Postings

	DV: Posted AI Job Out-of-State (1)	DV: Nr AI Job Out- of-State (logged) (2)	DV: AI Job Flight (3)
Proposed	0.0077* (0.0042)	0.3169** (0.1401)	0.0207* (0.0116)
Time-varying state controls	Yes	Yes	Yes
Company FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Observations	16,669	241	2,247
Adjusted R ²	0.2982	0.4688	0.5424

Panel B: Heterogeneous Effect of Regulation on Out-of-state AI Job Postings

	DV: Posted AI Job Out-of-State (4)	DV: Nr AI Job Out- of-State (logged) (5)	DV: AI Job Flight (6)
Proposed#Procedural Regulations	-0.0060 (0.0067)	0.0410 (0.3683)	-0.0015 (0.0145)
Proposed#Constraining Regulations	0.0103* (0.0061)	0.2459 (0.2183)	0.0221 (0.0176)
Time-varying state controls	Yes	Yes	Yes
Company FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Age FE	Yes	Yes	Yes
Observations	16,669	241	2,247
Adjusted R ²	0.2982	0.4624	0.5421

Appendix I: Regulatory Exposure Measure

In this Appendix I show some examples of the regulatory exposure measure using the cosine similarity score between the startup description and regulatory summary.

Consider the following company description: “Our mission at AQX is to bridge the gap between humans and machines, making AI audio communication universally accessible and strikingly natural. Unlock the future of business contact centers, cuttingedge technology meets the warmth of human conversation. Say goodbye to robotic interactions, experience seamless, intelligent communication that sets a new standard eliminating employee turnover and working at x lower operating costs.”

The two regulations with the highest similarity scores are as follows:

“Introduced on October 27, 2023, A8195, the Advanced Artificial Intelligence Licensing Act, requires the registration and licensing of high-risk advanced artificial intelligence systems, establishes the advanced artificial intelligence ethical code of conduct, and prohibits the development and operation of certain artificial intelligence systems.” - similarity score of 0.7895

“Introduced on January 4, 2023, A216, would require advertisements to disclose the use of synthetic media. Synthetic media is defined as “a computer-generated voice, photograph, image, or likeness created or modified through the use of artificial intelligence and intended to produce or reproduce a human voice, photograph, image, or likeness, or a video created or modified through an artificial intelligence algorithm that is created to produce or reproduce a human likeness.” Violators would be subject to a \$1,000 civil penalty for a first violation and a \$5,000 penalty for any subsequent violation.” - similarity score of 0.7856

Both regulations, upon reading, are highly relevant to the company’s description. The company clearly targets AI audio communication, highlighting the strikingly natural communication, which could be mistaken for human conversation. Next, I consider the regulations with the lowest similarity score:

“Introduced on November 3, 2023, S7735 (assembly version A7906), provides that it shall be unlawful for a landlord to implement or use an automated decision tool, unless it: (1) no less than annually, conducts a disparate impact analysis to assess the actual impact of any automated decision tool and publicly files the assessment; and (2) notifies all applicants that an automated decision tool will be used and provides the applicant with certain disclosures related to the automated decision tool. If passed, the law will go into immediate effect.” - similarity score 0.7360

“Restricts the use by an employer or an employment agency of electronic monitoring or an automated employment decision tool to screen a candidate or employee for an employment decision unless such tool has been the subject of a bias audit within the last year and the results of such audit have been made public; requires notice to employment candidates of the use of such tools; provides remedies; makes a conforming change to the civil rights law.” - similarity score 0.7337

Both these regulations are very particular contexts, which seem to be unrelated to the description of the company. Hence, the similarity score measure seems to have accurately ranked the most and least relevant AI regulations to the company in question.

Appendix II: AI Policy Enactments

In this appendix I discuss the methodology and results for the effect of AI policy enactments on fundraising outcomes.

A. Methodology of estimating the effect of policy enactments

There is usually a gap in time between regulation proposals and enactments (277 days on average). Once an AI regulation is proposed in the state chambers, it can quickly become a matter of public debate.¹³ Different interest groups, such as industry participants, may begin to lobby state governments to discourage legislation implementation that they deem harmful (Wei et al., 2024).¹⁴ I therefore adopt an Instrumental Variable (IV) approach to study the effect of enacting AI regulation.

Using a methodology similar in spirit to Acemoglu et al. (2022b), I restrict the sample to startups in states where AI policy proposals have been introduced, taking the timing of these proposals as given. To instrument for the likelihood of a state enacting AI regulation, I use the historical partisan composition of the state's senate, with data sourced from Ballotpedia (2024). More specifically, I employ the average proportion of Democratic senate members over the six years prior to the first AI regulation proposal.¹⁵ Since Democrats have traditionally been more

¹³ One such example is the recent SB 1047, which sparked significant criticism from academics, politicians, and investors, who expressed concerns about potential entry barriers and restrictions on innovation (see Hammond and Criddle 2024a; Hammond and Criddle 2024b).

¹⁴ This could lead to anticipation effects. In fact, attempting to run model (2) comparing startups in states with enacted policies to those in states without enactments gives rise to pre-trends for some of the specifications. Hence, staggered difference-in-difference models are not suitable in the context of regulation enactments.

¹⁵ Out of the 28 states with proposed AI regulation, 12 have a 2-year election cycle for their senate members. For these states, my instrument averages the Democratic proportion across three election terms. The remaining states follow a 4-year election cycle, except for Washington, D.C., which does not have a state senate due to its unique

regulation-friendly (Werthman, 2018), I expect states with more Democratic senates to be more likely to enact AI regulation, satisfying the relevance condition.

The exclusion restriction requires that the past Democratic composition of the senate does not affect startup-level outcomes directly, other than through the enactment of AI regulation. Specifically, it should not correlate with lobbying pressures after regulation is proposed. One concern is that lobbying pressures in general might influence the partisan composition of the senate across election cycles. However, this risk is mitigated by using the long-run moving average of partisan composition prior to regulation proposals, which reduces the influence of short-term political shifts.

Another concern with this instrument is that Democratic-majority states may generally be less business-friendly. Startup fixed effects mitigate this risk by controlling for any time-invariant characteristics specific to each startup, such as how they respond to broader regulatory environments. This means that persistent factors, like the overall political stance of Democratic-led states, are absorbed. Additionally, I include state-senate-election year fixed effects to account for political cycles and election-specific shifts that could affect both regulation and business performance, ensuring the observed effects are more likely due to AI regulation itself. The second-stage equation to be estimated is the following:

$$y_{ist} = \beta^{en} Enactment_s \times Post_{st} + ElectionYr_{st} + Age_{ist}^{en} + \delta^{en} Z_{st-1} + \alpha_{is}^{en} + \alpha_t^{en} + \xi_{ist} \quad (3)$$

status. Since Washington, D.C. is not technically a state, it is excluded from the regressions as the instrument is not applicable.

where $Enactment_s$ is a dummy variable about whether state s enacted an AI regulation, $Post_{st}$ is a dummy variable equal to 1 in time periods after the first AI regulation was proposed and $ElectionYr_{st}$ are a set of state-senate-election year dummies. The rest of the controls are the same as in equation (1) and (2). Equation (3) compares startups in states with policy enactments to startups in states with policy proposals, but no enactments. Thus, β^{en} measures the average change in startup performance attributable to the enactment of AI regulation. Letting $DemShare_s$ denote the average share of democrats in state's s senate in the 6-year period prior to its first proposal, the first-stage equation is as follows:

$$Enactment_s \times Post_{st} = DemShare_s \times Post_{st} + \alpha_{is}^{en} + \alpha_t^{en} + Age_{ist}^{en} + \delta^{en} Z_{st-1} + \xi_{ist} \quad (4)$$

B. Empirical results

AI regulation proposals act similarly to announcement effects, generating initial uncertainty and concern about potential future regulations, which can influence firms' behavior before the regulations are even enacted. To further explore this dynamic, I also investigate the effects of regulation enactments, focusing on states that have already introduced AI regulations. By doing so, I analyze whether the formal implementation of regulations reinforces or escalates the initial uncertainty caused by the proposals, giving insight into how the progression from proposal to enactment shapes the regulatory impact on businesses over time.

Appendix Table A.V shows OLS and 2SLS estimates, following the methodology described in the previous section. Regressions (1) and (3) display OLS estimates from the uninstrumented equation (3), while regressions (2) and (4) show 2SLS estimates, where enactments are instrumented by the 6-year average proportion of Democratic senate members prior to the first AI regulation proposal. All regressions are limited to startups in states with proposed regulation.

Insert Table A.V about here


Focusing on columns (1) and (2), which examine the determinants of the probability of fundraising success, the OLS coefficient for the ‘Proposed’ dummy is smaller in absolute value compared to the 2SLS estimate. This suggests an upward bias in the OLS results, consistent with the idea that lobbying efforts lead to the enactment of less harmful regulations. According to the 2SLS estimates, the enactment of regulations leads to a 3% points annual decrease in the probability of fundraising, which is larger in magnitude than the 2% points decrease associated with regulation proposals.

Appendix III: Additional Tables and Figures


Figure A.I: Dynamic Effect of AI regulation proposals on firm governance hiring

This figure illustrates the dynamic effects of AI regulation proposals on startups' governance hiring, as estimated by equation (2) in Section II.A. The dots represent point estimates of the dynamic difference-in-differences coefficients, while the bars show the 95% confidence intervals. The year prior to the regulation proposal is normalized to zero. In Panel A, the dependent variable is a binary indicator for whether a startup posted an AI governance position. In Panel B, the dependent variable is the natural logarithm of the number of AI governance positions posted, conditional on having posted at least one such position. In Panel C, the dependent variable is the ratio of governance to non-governance AI Jobs posted in the particular year. The model controls for the lagged log level and growth rate of real GDP, lagged log disposable income per capita, startup fixed effects, year fixed effects, and age fixed effects. Standard errors are clustered at the state level.

Panel A: Posted Governance Job



Panel B: Nr. Governance Jobs Posted (logged)



Panel C: Fraction Governance/ Non-Governance AI Jobs






Figure A.2: Dynamic Effect of AI regulation proposals on firm Applied & Frontier AI hiring

This figure illustrates the dynamic effects of AI regulation proposals on startups' Applied & Frontier AI hiring, as estimated by equation (2) in Section II.A. The dots represent point estimates of the dynamic difference-in-differences coefficients, while the bars show the 95% confidence intervals. The year prior to the regulation proposal is normalized to zero. In Panel A (C), the dependent variable is a binary indicator for whether a startup posted an Applied (Frontier) AI position. In Panel B (D), the dependent variable is the natural logarithm of the number of Applied (Frontier) AI positions posted, conditional on having posted at least one such position. Panel E depicts the dynamic effects of regulation on the ratio between Frontier and Applied AI Jobs posted. The model controls for the lagged log level and growth rate of real GDP, lagged log disposable income per capita, startup fixed effects, year fixed effects, and age fixed effects. Standard errors are clustered at the state level.


Panel A: Posted Applied AI Job




Panel B: Nr. Applied AI Jobs Posted (logged)



Panel C: Posted Frontier AI Job



Panel D: Nr. Frontier AI Jobs Posted (logged)



Panel E: Ratio Frontier - Applied AI Jobs






Figure A.3: Dynamic Effect of AI regulation proposals on firm out-of-state hiring

This figure illustrates the dynamic effects of AI regulation proposals on startups' hiring in non-regulated states, as estimated by equation (2) in Section II.A. The dots represent point estimates of the dynamic difference-in-differences coefficients, while the bars show the 95% confidence intervals. The year prior to the regulation proposal is normalized to zero. In Panel A, the dependent variable is a binary indicator for whether a startup posted an AI job (Applied or Frontier AI) in a non-regulated state as of the end of 2023, which is not the state where the startup is headquartered. In Panel B, the dependent variable is the natural logarithm of the number of AI positions posted in non-regulated states, conditional on having posted at least one such position. In Panel C, the dependent variable is the ratio of the AI jobs posted in non-regulated states, over the total AI jobs posted in a particular year. The model controls for the lagged log level and growth rate of real GDP, lagged log disposable income per capita, startup fixed effects, year fixed effects, and age fixed effects. Standard errors are clustered at the state level.

Panel A: Posted AI Job Out of State



Panel B: Nr. AI Jobs posted Out of State (logged)



Panel C: AI Job flight




Table A.I Alternative Estimation Methods for the Effect of AI Regulation Proposals

This table presents the results from two alternative estimation methods of equation (1) in Section II.A, evaluating the effect of AI regulation proposals on startups' fundraising performance. Columns (1) and (2) report results using the Sun & Abraham (2021) (SA DiD) estimator, whereas columns (3) and (4) report results using the Callaway and Sant'Anna (2021) (CS DiD) estimator. The control group is the never-treated startups for all regressions (1)-(4). The dependent variable in column (1) and (3) is a binary indicator for whether a startup raised funding in a given year. The dependent variable in column (2) and (4) is the natural logarithm of the amount raised, conditional on successful fundraising. Time-varying state controls refer to the lagged level (logged) and growth rate of Real GDP as well as lagged log of disposable income per capita. Due to multi-dimensional fixed effects, 145 singleton observations are dropped in column (1), and 2,362 are dropped in column (2). The CS DiD estimator requires common support between treated and untreated units, meaning that some units are excluded if their covariates do not sufficiently overlap. To address this, I modified the controls to improve comparability. Specifically, I replaced age dummies with age and age-squared, and standardized the log of Real GDP (removing the state mean and dividing by the state's standard deviation) to align treated and control units. These adjustments increase common support and reduce the number of excluded observations. The same transformations applied to the TWFE and SA DiD estimators yield similar results as the original. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

	SA DiD Estimator		CS DiD Estimator	
	(1) Raised Eq	(2) Amount Raised (log)	(3) Raised Eq	(4) Amount Raised (log)
Proposed	-0.0520*** (0.0068)	0.0653 (0.0913)	-0.0579*** (0.0215)	0.3724 (0.8508)
Time-varying state control	Yes	Yes	Yes*	Yes*
Company FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes*	Yes*
Observations	99,729	9,367	91,839	4,039
Adjusted R ²	0.2492	0.7103	.	.

Table A.II Robustness Check Excluding 2022&2023 – Effect of AI Regulation Proposals on Fundraising Performance

This table presents the Two-Way Fixed Effects (TWFE) difference-in-difference estimates based on equation (1) from Section II.A, assessing the effect of AI regulation proposals on startups' fundraising performance. As a robustness check, the analysis excludes observations from 2023 in Panel A, since a large number of regulations were proposed in this year. Panel B additionally excludes observations from 2022. In columns (1) and (3), the dependent variable is a binary indicator for whether a startup raised funding in a given year. In column (2) and (4), the dependent variable is the natural logarithm of the amount raised, conditional on successful fundraising. Time-varying state controls refer to the lagged level (logged) and growth rate of Real GDP, as well as lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

Panel A: Effect of Policy Proposal when excluding 2023		
	TWFE DiD Estimator	
	Raised Eq	Amount Raised (log)
	(1)	(2)
Proposed	-0.0289*** (0.0064)	0.1039 (0.0626)
Time-varying state controls	Yes	Yes
Company FE	Yes	Yes
Year FE	Yes	Yes
Age FE	Yes	Yes
Observations	85,935	8,587
Adjusted R ²	0.2589	0.7130
Panel B: Effect of Policy Proposal when excluding 2022&2023		
	(3)	(4)
Proposed	-0.0276*** (0.0051)	0.0838 (0.0503)
Time-varying state controls	Yes	Yes
Company FE	Yes	Yes
Year FE	Yes	Yes
Age FE	Yes	Yes
Observations	72,412	7,446
Adjusted R ²	0.2649	0.7118

Table A.III Robustness Check Placebo Timing – Effect of AI Regulation Proposals on Fundraising Performance

This table presents the Two-Way Fixed Effects (TWFE) difference-in-difference estimates from equation (1) in Section II.A, assessing the effect of AI regulation proposals on startups' fundraising performance. As a robustness check, placebo timings are used for the first AI regulation in each state. Panel A shifts the timing back by 2 years, while Panels B and C shift it by 4 and 6 years, respectively. In columns (1), (3), and (5), the dependent variable is a binary indicator for whether a startup raised funding. In columns (2), (4), and (6), it is the natural log of the amount raised, conditional on successful fundraising. Treated startups are defined as those experiencing a shift in "Placebo Regulation" proposals. As the timing shifts further back, the sample size decreases, with only older startups remaining. Results are consistent when using the baseline sample size from Table I. Controls include lagged real GDP (logged), its growth rate, disposable income per capita (logged), startup, year, and age fixed effects. Standard errors are clustered at the state level, with significance indicated by ***, **, and * for 99%, 95%, and 90% confidence levels.

Panel A: Placebo timing - 2 years back		
	TWFE DiD Estimator	
	Raised Eq	Amount Raised (log)
	(1)	(2)
Proposed	-0.0012 (0.0087)	0.0476 (0.0668)
Controls	Yes	Yes
Observations	81,938	7,484
Adjusted R ²	0.2495	0.7117
Panel B: Placebo Timing - 4 years back		
	(3)	(4)
Proposed	0.0055 (0.0086)	0.1375* (0.0780)
Controls	Yes	Yes
Observations	60,144	5,269
Adjusted R ²	0.2561	0.7071
Panel B: Placebo Timing - 6 years back		
	(5)	(6)
Proposed	0.0036 (0.0129)	0.0673 (0.1525)
Controls	Yes	Yes
Observations	38,484	3,055
Adjusted R ²	0.2556	0.6829

**Table A.IV: Robustness Check Excluding each state in turn - Effect of AI Regulation
Proposals on Probability of financing**

This table reports two-way fixed effects (TWFE) difference-in-differences (DiD) estimates from equation (1) in Section II.A, measuring the effect of AI regulation proposals on startups' fundraising performance. To assess robustness, I re-estimate the model excluding one state at a time. The dependent variable is a dummy for whether a startup raised any funding. The specification includes the log of lagged real GDP, real GDP growth, and the log of disposable income per capita, along with startup, year, and age fixed effects. Standard errors are clustered at the state level, with ***, **, and * denoting significance at the 1%, 5%, and 10% levels.

	Coefficient	Standard Error	p-value
All states (Baseline)	-0.0231***	0.0056	0.0002
Exclude CA	-0.0159*	0.0083	0.0618
Exclude CO	-0.0251***	0.0056	0.0000
Exclude CT	-0.0232***	0.0057	0.0002
Exclude DC	-0.0226***	0.0057	0.0002
Exclude DE	-0.0229***	0.0057	0.0002
Exclude GA	-0.0228***	0.0058	0.0003
Exclude HI	-0.0232***	0.0056	0.0001
Exclude IL	-0.0251***	0.0055	0.0000
Exclude IN	-0.0230***	0.0057	0.0002
Exclude MA	-0.0223***	0.0058	0.0003
Exclude MD	-0.0240***	0.0057	0.0001
Exclude ME	-0.0230***	0.0056	0.0002
Exclude MN	-0.0229***	0.0057	0.0002
Exclude MT	-0.0230***	0.0057	0.0002
Exclude NH	-0.0230***	0.0056	0.0002
Exclude NJ	-0.0233***	0.0058	0.0002
Exclude NM	-0.0231***	0.0057	0.0002
Exclude NY	-0.0213***	0.0071	0.0045
Exclude OR	-0.0235***	0.0057	0.0001
Exclude PA	-0.0227***	0.0057	0.0002
Exclude RI	-0.0230***	0.0056	0.0002
Exclude SC	-0.0231***	0.0057	0.0002
Exclude TN	-0.0232***	0.0057	0.0002
Exclude TX	-0.0238***	0.0063	0.0004
Exclude VA	-0.0247***	0.0056	0.0001
Exclude VT	-0.0230***	0.0056	0.0002
Exclude WA	-0.0225***	0.0062	0.0007
Exclude WV	-0.0231***	0.0056	0.0002

Table A.V: Robustness Check for AI regulation enactments - 2SLS results

This table presents OLS and 2SLS regression estimates for the effect of AI regulation enactments on startup financing outcomes, for startups in states with proposed regulation. Regressions (1) and (3) display estimates from the uninstrumented equation (3), while regressions (2) and (4) show 2SLS estimates, where enactments are instrumented by the 6-year average proportion of Democratic senate members, as described in Appendix II. All regressions are limited to startups in states with regulation proposals. The dependent variable in columns (1) and (2) is a binary indicator for whether a startup raised funding in a given year. In columns (3) and (4), it is the logged amount raised, conditional on fundraising success. Time-varying state controls include the lagged (logged) level and growth rate of real GDP and lagged (logged) disposable income per capita. Standard errors are clustered at the state level, with significance levels denoted by ***, **, and * for 99%, 95%, and 90% confidence, respectively.

	DV: Rasied Eq		DV: Amount Raised (log)	
	OLS (1)	IV Estimate (2)	OLS (3)	IV Estimate (4)
Enacted	-0.0138* (0.0076)	-0.0316*** (0.0104)	0.0938 (0.0628)	0.1271** (0.0591)
Time-varying state controls	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Election Year FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	84,928	83,695	8,331	8,260
Adjusted R ² (Centered for IV)	0.2439	0.0014	0.7095	0.0026
First-Stage for IV Regressions				
DemShare		1.2418*** (0.2024)		1.3764*** (0.1315)
F-Stat		37.6297		109.5393

Table A.VI: AI regulation taxonomy

This table catalogs operative clauses and definitions for procedural and constraining rules, reports their distribution across regulations and states, and includes brief excerpts illustrating each provision in practice.

Panel A: Procedural Regulations

Mechanism	Definition	Justification for inclusion under the regulation type.	Example	Nr of Regulations	Nr of States
				<i>Total: 56 regulations across 23 states.</i>	
Impact Assessments	Must conduct a formal risk or impact assessment before deployment (1/0)	Included as procedural regulation. An impact assessment obliges providers to conduct a formal, documented evaluation of reasonably foreseeable risks and harms before launch and at regular intervals, turning engineering artifacts into mitigation commitments. The example requires companies to perform regular risk assessments across design, development, and implementation to identify, assess, and mitigate foreseeable risks and cognizable harms. This ex ante and ongoing duty creates a verifiable record for audits and disclosures and links deployment to concrete monitoring and remediation plans.	... "The company shall conduct <i>regular risk assessments</i> to identify, assess and mitigate reasonably foreseeable risks and cognizable harms related to their products and services, including in the design, development and implementation of such products and services."... (H. 6286, Rhode Island General Assembly, 2023, p.2).	47	23
Disclosure to regulators	Must disclose technical details to regulator (1/0)	Included as procedural regulation. Disclosures to regulators convert internal engineering knowledge into supervisory visibility. The example requires an employer to submit a system summary - describing the monitoring method, the number of affected workers, the data to be collected, and how that data will inform employment decisions. Such targeted disclosures let agencies assess proportional risk, set conditions or halt deployment if needed, establish a documentary baseline for later audits and investigations, and tie specific systems to accountable parties without requiring full public release.	... "Before an employer uses an electronic productivity system, the employer <i>shall submit a summary of the system to the department</i> , including information on the specific form of monitoring, the number of workers impacted, the data that will be collected, and how that data will be used in making employment-related decisions."... (H. 1873, Massachusetts General Court, 2023, p.19).	50	22
Third-party audits	Requires third-party audit, certification, or testing (Alanoca et al., 2025) (1/0)	Included as procedural regulation. Third-party audits, certifications, or pre-market testing introduce independent verification of safety, performance, and compliance, converting internal claims into enforceable attestations and gating deployment. In the example provided, a hospital may use a diagnostic algorithm only after confirming it has been certified by the Departments of Public Health and Innovation and Technology, creating an ex ante check by expert authorities. This mechanism standardizes evaluation criteria, produces an auditable record for regulators, conditions procurement on meeting verified thresholds, and deters deployment of unvalidated systems.	... "Diagnostic algorithm. Before using any diagnostic algorithm to diagnose a patient, the University of Illinois Hospital must first confirm all of the following: (1) <i>The diagnostic algorithm has been certified</i> by the Department of Public Health and the Department of Innovation and Technology."... (H.B. 1002, Illinois General Assembly, 2023, p.2).	9	6
Registration of AI tools	Must register or license AI tools (1/0)	Included as procedural regulation. Registration requirements create an authoritative inventory of systems and operators, tying specific models to accountable entities as a precondition of operation. The example - requiring any company operating a large-scale generative model to register with the attorney general within 90 days - establishes a time-bound filing that identifies the operator, model, and intended uses, supplies a supervisory contact point, and enables oversight to scale. By making registration a gating duty, the rule deters fly-by-night deployments, supports incident response and recall, and links subsequent obligations (audits, disclosures, enforcement) to a verified record.	... "Any company operating a large-scale generative artificial intelligence model <i>shall register with the attorney general</i> within 90 days of the effective date of this act."... (S.31, Massachusetts General Court, 2023, p. 3).	7	6

Panel B: Constraining Regulations

Mechanism	Definition	Justification for inclusion under the regulation type.	Example	Nr of Regulations	Nr of States
				<i>Total: 37 regulations across 14 states.</i>	
Broad AI Definition	Regulation defines "artificial intelligence" in a way that encompasses a wide range of technologies, including both rule-based systems (e.g., expert systems, logic-based decision trees) and learning-based systems (e.g., machine learning, neural networks). (1/0)	Included as constraining regulation. A broad, technology-neutral definition of "AI system" functions as the coverage gate, determining who is in scope before risk thresholds, prohibitions, and emergency brakes apply. In the cited example, AI is defined in vague, functional terms ("systems able to perform tasks that normally require human intelligence"), which broadens inclusion and makes it more likely that a company's products and operations would be covered.	... "As used in this section, the term "artificial intelligence" shall mean the theory and development of computer systems able to perform tasks <i>that normally require human intelligence</i> .".... (H.B. 49, Pennsylvania General Assembly, 2023, p. 3).	30	13
High-Risk Category	There is a defined category of "high-risk" AI (1/0)	Included as constraining regulation. A statutory "high-risk" category draws the boundary at risk thresholds, determining when heightened duties attach. Definitions tied to function and use—such as systems used, or reasonably foreseeable to be used, to make consequential decisions or to categorize by protected characteristics—broaden inclusion by capturing both present and near-term deployments, regardless of model size or supervision structure. The foreseeability standard closes timing and relabeling gaps, anchoring coverage where harms concentrate, and thereby narrowing the feasible set of applications ex ante while making obligations attach predictably to systems with the greatest potential for harm.	..." <i>High-risk artificial intelligence system</i> " means any artificial intelligence system, regardless of the number of parameters and supervision structure, that is: (A) used, reasonably foreseeable as being used, or is a controlling factor in making a consequential decision; (B) used, or reasonably foreseeable as being used, to categorize groups of persons by sensitive and protected characteristics, such as race, ethnic origin, or religious belief; "... (H. 7111, Vermont General Assembly, 2023, p. 5).	3	2
Prohibited Use Cases	Contains full bans on some use-cases of AI (e.g., facial recognition) (1/0)	Included as constraining regulation. Categorical prohibitions on specified AI uses set bright-line limits ex ante, removing reliance on post hoc risk assessments or discretionary mitigation. The example imposes a full ban on incorporating facial, gait, or emotion recognition in electronic monitoring and automated decision systems, thereby eliminating those design and procurement options from the feasible set. Such rules preclude "paper compliance" and simplify enforcement through clear, observable predicates.	..." <i>Prohibition on facial, gait, and emotion recognition technology</i> . Electronic monitoring and automated decision systems shall not incorporate any form of facial, gait, or emotion recognition technology." ... (H. 114, Vermont General Assembly, 2023, p.12).	13	8
Temporary Injunctions	Legal orders that temporarily halt the development, deployment, or continued operation of AI systems suspected of posing significant risks or violating regulations (Alanoca et al., 2025, p. 5). (1/0)	Included as constraining regulation. Temporary injunction authority functions as an emergency brake: it empowers enforcers to seek a court order halting the deployment or continued operation of AI systems suspected of violating statutory duties or posing significant risks. The example authorizes the Attorney General to pursue temporary or permanent injunctions upon a "reason to believe" showing, enabling rapid containment before harms become irreversible. This tool narrows the feasible set ex ante by deterring risky launches and forcing prompt remediation, while preserving due process through judicial oversight.	... "In any case in which the Attorney General for the District of Columbia has reason to believe that any person has used, is using, or intends to use any method, act, or practice in violation of this act or a regulation promulgated under this act, or has failed to provide a notice, a disclosure, or a report required by this act, the Attorney General for the District of Columbia may commence appropriate civil action in the Superior Court of the District of Columbia for: (1) <i>A temporary or permanent injunction</i> ;" ... (B. 25-0114, Council of the District of Columbia, 2023, p.14).	13	7
Open Source	Covers open-source models and libraries (1/0)	Included as constraining regulation. An explicit rule covering open-source models and libraries fixes the boundary of who is in scope in open ecosystems. The example's formulation—defining "open code" as publicly available code that the public can use, modify, or distribute "irrespective of ... the underlying license agreement"—broadens inclusion, reduces ambiguity, limits opportunities for regulatory arbitrage (e.g., offloading risky components by releasing them "as open"), and ensures that obligations attach predictably to both upstream developers and downstream deployers, thereby shaping design and distribution choices ex ante.	... " <i>open code</i> " shall mean software whose code is made available to the public where the public has the ability to use, modify or distribute the code <i>irrespective</i> of the associated costs to use, modify or distribute the code, if any, or <i>the underlying license agreement</i> "... (A. 8105, New York State Assembly, 2023, p.2).	2	2

Table A.VII: Robustness Check for the heterogenous effect of AI regulation - dynamic version of procedural and constraining regulation

This table reports difference-in-differences estimates of the heterogeneous effects of AI regulation proposals on startup fundraising and exit. “Proposed” equals 1 for state–year observations beginning in the first year a state proposes an AI regulation. “Procedural Regulations (dynamic)” and “Constraining Regulations (dynamic)” equal 1 from the first year a state proposes the respective rule and remain 1 thereafter. The dependent variable in columns (1) - (3) is a binary indicator of whether a startup has raised funding. The dependent variable in columns (4) - (6) is a binary indicator for whether the startup was acquired. Startups exit the panel upon acquisition; for columns (4)–(6), the sample retains only treated startups whose acquisition occurs after the proposal date. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

	DV: Raised Eq			DV: Acquired		
	(1)	(2)	(3)	(5)	(6)	(7)
Proposed	-0.0336*** (0.0100)	0.0080 (0.0087)		0.0180*** (0.0023)	0.0093*** (0.0025)	
Procedural Regulations (Dynamic)	0.0196 (0.0166)		0.0201* (0.0110)	-0.0064** (0.0030)		-0.0020 (0.0028)
Constraining Regulations (Dynamic)		-0.0388*** (0.0097)	-0.0416*** (0.0064)		0.0065** (0.0031)	0.0154*** (0.0027)
Time-varying state controls	Yes	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	99,729	99,729	99,729	94,439	94,439	94,439
Adjusted R ²	0.2480	0.2482	0.2483	0.0530	0.0530	0.0527

Table A.VIII: Definitions and Summary Statistics on AI hiring

This table presents definitions (Panel A) and summary statistics (Panel B) on AI hiring, covering 2,883 startups matched to Lightcast's job postings.

Panel A: Variable Definitions

Posted Governance Job	A dummy variable equal to 1 if the company has posted at least one AI governance job each calendar year since incorporation. Source: Lightcast jobs postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Posted Applied AI Job	A dummy variable equal to 1 if the company has posted at least one Applied AI job each calendar year since incorporation. Source: Lightcast jobs postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Posted Frontier AI Job	A dummy variable equal to 1 if the company has posted at least one Frontier AI job each calendar year since incorporation. Source: Lightcast jobs postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Posted AI Job Out of State	A dummy variable equal to 1 if the company has posted at least one AI job (Frontier or Applied AI) each calendar year since incorporation in a non-regulated states as of 2023, other than the one it is headquartered. Source: Lightcast jobs postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Nr. Governance Jobs	Count of AI governance jobs posted by the company in each calendar year since incorporation, conditional on posting. Source: Lightcast job postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Nr. Applied AI Jobs	Count of Applied AI jobs posted by the company in each calendar year since incorporation, conditional on posting. Source: Lightcast job postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Nr. Frontier AI Jobs	Count of Frontier AI jobs posted by the company in each calendar year since incorporation, conditional on posting. Source: Lightcast job postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Nr. AI Jobs Out of State	Count of AI jobs (Frontier or Applied AI) posted by the company each calendar year since incorporation in a non-regulated states as of 2023, other than the one it is headquartered. Source: Lightcast jobs postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
DV: Ratio Governance/Non-Gov AI Jobs	Ratio of AI-governance jobs to Non-Governance AI jobs posted by the company in a given calendar year. Source: Lightcast job postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
Ratio Frontier AI/ Applied AI	Ratio of Frontier AI (Exploration) jobs to Applied AI (Exploitation) jobs posted by the company in a given calendar year. Source: Lightcast job postings classified according to the definitions in Appendix IV and the methodology in Appendix V.
AI Job Flight	Ratio of AI jobs in non-regulated states (other than the headquarter state) to the total AI jobs posted by the company in a given calendar year. Source: Lightcast job postings classified according to the definitions in Appendix IV and the methodology in Appendix V.

Panel B: Summary Statistics

	Overall Frequency of Job Types across companies					
	Total Nr. Job Postings			Total Nr. Companies that posted job		
Governance Job	6,661			554		
Applied AI Job	13,847			1,130		
Frontier AI Job	3,735			562		
AI Job Out of State (Frontier or Applied)	1,779			188		
	Summary Statistics at the Company-Year level					
	Obs.	Mean	St. Dev.	p5	p50	p95
Posted Governance Job	16,717	0.0786	0.2691	0.0000	0.0000	1.0000
Posted Applied AI Job	16,717	0.1648	0.3711	0.0000	0.0000	1.0000
Posted Frontier AI Job	16,717	0.0680	0.2517	0.0000	0.0000	1.0000
Posted AI Job Out of State	16,717	0.0211	0.1436	0.0000	0.0000	0.0000
Nr. Governance Jobs	1,314	5.0693	12.5108	1.0000	2.0000	17.0000
Nr. Applied AI Jobs	2,756	5.0243	12.8656	1.0000	2.0000	17.0000
Nr. Frontier AI Jobs	1,136	3.2879	6.3831	1.0000	1.0000	10.0000
Nr. AI Jobs Out of State	352	5.0540	11.0330	1.0000	2.0000	21.0000
Ratio Governance/Non-Governance	3,100	0.3042	1.3354	0.0000	0.0000	1.5000
Ratio Frontier AI/Applied AI	2,756	0.2534	0.8471	0.0000	0.0000	1.0000
AI Job Flight	3,122	0.0564	0.1978	0.0000	0.0000	0.5000

Table A.IX: AI regulation and strategic pivoting

This table reports difference-in-differences estimates of the effect of AI regulation proposals on startup Applied & Frontier AI hiring. “Proposed” equals 1 from the first year a state proposes an AI regulation. Dependent variables are: any Applied (Frontier) AI posting in column 1 (3), the number of Applied (Frontier) AI postings conditional on at least one such posting in column 2 (4) and the ratio between Frontier and Applied AI postings in column 5. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with the lagged log of disposable income per capita. Standard errors are clustered at the state level. Significance: ***, **, * denote 99%, 95%, and 90% confidence levels.

	DV: Posted Applied AI Job (1)	DV: Cond. Nr. Applied AI Jobs (logged) (2)	DV: Posted Frontier AI Job (3)	DV: Cond. Nr Frontier AI Jobs (logged) (4)	DV: Ratio Frontier/Applied AI Job (5)
Proposed	0.0084 (0.0068)	-0.0330 (0.0485)	0.0088 (0.0069)	-0.1197 (0.1018)	0.0942 (0.1020)
Time-varying state controls	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes
Observations	16,669	2,270	16,669	831	2,270
Adjusted R ²	0.3361	0.4255	0.3154	0.425	0.1834

Appendix IV: Classification taxonomy for AI labor demand

This appendix describes the task-based taxonomy used to classify AI job postings into mutually exclusive categories—**AI Exploitation**, **AI Exploration**, and **Non-AI**—along with an orthogonal **AI-Governance** indicator. The taxonomy is designed to capture firms’ intended marginal capability acquisition at the time reallocation decisions are made, consistent with the exploration–exploitation framework outlined in Section IV. It maps postings onto the conceptual margins along which regulation and financing are expected to shift technical effort and forms the basis for all empirical tests of reallocation across exploration, exploitation, and governance.

The unit of observation is the individual job posting. Each posting is assigned a single primary category according to its dominant task profile and core deliverables. Classification draws on the full job text—including the title, listed responsibilities, required skills, and any performance or accountability metrics. When signals are mixed or ambiguous, the tie-breaker rules described in Section A are applied.

A. Primary Categories

Postings are classified into three mutually exclusive categories based on their dominant task profile and core deliverables: AI Exploitation (Applied AI), AI Exploration (Frontier AI) and Non-AI. If a posting satisfies neither set of criteria, it is coded as Non-AI (e.g., AI-adjacent infrastructure or roles without ML/AI accountability).

A.1. AI Exploitation (Applied AI)

Definition.

Roles whose primary output is the refinement, scaling, and operation of existing ML/AI capabilities using established methods. Core deliverables include building and maintaining ML pipelines, serving/inference, monitoring and drift detection, reliability and latency/throughput targets, safe rollouts, and iteration to product key performance indicators (KPIs). These roles do not own invention of new learning algorithms as a core deliverable.

Inclusion criteria.

Classify a posting as AI Exploitation when teams are accountable for:

- (i) model behavior in production (quality, safety, product ML metrics), or
- (ii) ML-specific serving/training pipelines and their service level objectives (SLOs)—including inference latency and throughput, model freshness, drift monitoring, and rollout safety—regardless of ownership of model parameters.

Exclusion criteria/ Not–AI.

Postings whose deliverables are generic infrastructure service level objectives (for example, cluster uptime and cost, storage, or generic networking) and that carry no accountability for ML features, ML pipelines, or ML product metrics should be coded as Not–AI.

A.2. AI Exploration (Frontier AI)

Definition.

Roles whose primary output is search and experimentation to create new models/algorithms, or materially new training/optimization schemes. While deployment may occur, the defining deliverable is methodological innovation that extends the firm's technical frontier.

Subtypes.

Recorded as sub-labels (both map to Exploration):

- **Exploration—Applied:** method innovation directed at improving a product or use case.
- **Exploration—Fundamental:** core R&D producing new algorithms, theory, benchmarks, or publications.

Inclusion criteria (sufficient, not all required).

Positive indicators include explicit novelty deliverables (“new/novel model/algorithm,” research agenda, top-venue publication goals); algorithmic advances (e.g., active learning, simulation-driven design, custom objectives/losses, meta/curriculum learning); system–algorithm co-design that alters how learning or inference works (compiler/kernel/scheduling/quantization/distributed-training innovations); retrieval methodology invention (e.g., learned retrievers, new negative-mining/selection schemes); data-centric algorithmic methods (programmatic labeling frameworks; new selection/synthesis algorithms).

Exclusion criteria (default to Exploitation unless overridden by positive signals).

Exclude postings focused on pipeline building, model serving, monitoring, and service level objective (SLO)/cost work using established methods; fine-tuning or serving large language models (LLMs) and prompt engineering with standard recipes; feature development or user experience (UX) work without machine-learning (ML) method novelty; general infrastructure tasks, feature-store maintenance, or dataset wrangling without algorithmic innovation.

Tie-breakers and edge cases.

I classify postings as Exploration only when the text specifies credible and concrete novelty deliverables—such as developing a new model or algorithm, defining a new training objective, or designing a new retrieval method. Large-language-model fine-tuning that applies standard adapters or recipes to typical datasets is coded as Exploitation, as it implements established methods rather than generating new ones. Generic phrases such as “research and development,” “cutting-edge,” or “state-of-the-art” are not sufficient; the posting must identify a specific innovation output.

Managerial roles are coded according to their team's primary outputs: teams charged with producing new methods are Exploration, whereas those responsible for productionization, pipelines, and reliability are Exploitation. When descriptions are ambiguous or lack sufficient detail, I default to Exploitation as the more conservative assignment.

When a posting mixes operational and innovation responsibilities, it is classified as Exploration only if it explicitly includes credible novelty deliverables; otherwise, it is coded as Exploitation. This approach aligns the discrete label with the empirical object of interest—whether firms allocate any marginal effort to exploration. The presence of exploratory responsibilities satisfies that condition. It also corresponds to theoretical models of dynamic experimentation, in which firms allocate effort between a safe option and a risky, information-generating option (Bolton & Harris, 1999; Keller et al., 2005). Any positive allocation to the risky option constitutes exploration, so assigning mixed postings to Exploration provides the closest discrete analogue. Finally, this rule addresses a measurement concern: job advertisements often contain operational language even for research-oriented roles. Classifying mixed postings as Exploitation would undercount exploration and bias estimated reallocation effects toward zero.

B. AI-Governance (Orthogonal Indicator)

Definition.

An AI-Governance indicator (binary) is recorded in addition to the primary category (Exploitation or Exploration). The indicator equals one when a posting bears explicit responsibility for activities aimed at mitigating AI-related risks - such as documentation, audit, and oversight - and for ensuring compliance with AI- or data-governance laws, regulations, or standards.

Examples include roles referencing regulatory frameworks or standards such as the GDPR, CCPA/CPRA, HIPAA/FDA SaMD, the EU AI Act, SR 11-7 on model risk management, the NIST AI Risk Management Framework, ISO/IEC 42001, and ISO/IEC 23894.

Inclusion criteria.

Assign the AI-Governance indicator when core deliverables include formal risk assessments and controls for AI systems; documentation artifacts (model cards, data documentation, lineage/provenance records); conformity assessment, audit preparation, or attestation; approval or release gates tied to regulatory or policy requirements; privacy and security controls specific to machine-learning workflows; or ongoing compliance monitoring for bias, robustness, or safety metrics.

Exclusion criteria.

Do not assign the AI-Governance indicator when activities consist only of quality assurance, testing, red-teaming, or evaluation without ownership of AI regulatory or standards compliance; or when responsibilities are limited to general infrastructure or security with no machine-learning accountability.

Routing with primary categories.

The AI-Governance indicator does not create a separate category. Each posting is first classified as AI Exploitation or AI Exploration based on its dominant task profile, or as Non-AI if it satisfies neither set of criteria. The governance indicator is then recorded in parallel and coded as present whenever AI-governance ownership forms part of the role's remit. Examples of postings with AI-governance responsibilities include a legal or compliance role in Non-AI, an operations role with model-risk controls in Exploitation, and a research role focused on AI-safety measures and related assurance practices in Exploration.

Appendix V: Multi-Head Job Classification - Data and Methodology

This appendix describes a multi-task text classifier that assigns each job posting to one of three mutually exclusive categories—Applied AI, Frontier AI, or Non-AI—and, independently, flags a Governance indicator. The model fine-tunes a RoBERTa encoder that is shared across tasks and feeds two task-specific heads. The three-way head is trained with class-weighted cross-entropy; the Governance head is trained with a logistic loss on logits. The data, code, and split protocol are held fixed to ensure exact comparability and replication.

Data and label construction.

The annotated dataset comprises 3,132 job postings. Each observation includes the job text and labels for a three-way category—Applied AI, Frontier AI, or Non-AI—and a separate Governance flag. Text is tokenized to a maximum of 512 sub-tokens. To prevent order or indexing drift, the index is reset and a deterministic permutation over row positions is drawn with the seed set to 42. Stratification is applied on the joint label that combines the three-way category with the Governance flag so that the marginal distribution across tasks is preserved in each partition. The final split contains 2,192 observations for training, 470 for validation, and 470 for testing. Within the training split, class counts for the three-way task are 625 for Non-AI, 700 for Frontier AI, and 867 for Applied AI. Governance positives are relatively rare, with 144 in training, 31 in validation, and 31 in testing.

Model.

The classifier fine-tunes roberta-base with masked-mean pooling over the final hidden states. Two heads then project the pooled representation to logits: a two-layer perceptron for the three-class task and a two-layer perceptron for Governance. Linear layers use Xavier-uniform initialization with zeroed biases.

Training minimizes the sum of a class-weighted cross-entropy for the three-class head and a logistic loss on logits for the Governance head. The three-class loss uses label smoothing of 0.01 and class weights computed on the training split only—approximately 1.169 for Non-AI, 1.044 for Frontier AI, and 0.843 for Applied AI. The Governance loss applies a positive-class weight capped at 8 to address imbalance and is multiplied by a task weight of 0.3. Optimization uses AdamW with a learning rate of $2e-5$ and weight decay of 0.01, together with a linear schedule and a warmup ratio of 6%. The per-step batch size is 8 with gradient accumulation of 2, which yields an effective batch size of 16. cuDNN runs in deterministic mode with benchmarking disabled.

Checkpoints are selected by the macro-F1 score of the three-class task on the validation set. The Governance decision threshold is chosen on the validation set by sweeping the precision–recall curve and maximizing the F1 score, without imposing a minimum-precision constraint in this run. By convention, probabilities at or above the chosen threshold are mapped to the positive class. The tuned Governance threshold is 0.90.

Evaluation and test results.

All metrics are computed on the fixed test set of 470 observations using the single best checkpoint selected on validation. For the three-class task, the test macro-F1 score is 0.837 and overall accuracy is 0.832. Per-class test performance is as follows. For Non-AI, precision is 0.906, recall is 0.866, and F1 is 0.885 with support of 134. For Frontier AI, precision is 0.817, recall is 0.833, and F1 is 0.825 with support of 150. For Applied AI, precision is 0.794, recall is 0.806, and F1 is 0.800 with support of 186. For the Governance indicator, at the conventional probability cut-off of 0.5 the test precision is 0.806, the test recall is 0.935, and the test F1 is 0.866. Using the validation-tuned cut-off of 0.90, test precision is 0.848, test recall is 0.903, and test F1 is 0.875. Moving from the conventional cut-off to the tuned cut-off shifts the operating point toward higher precision with a modest reduction in recall, which is consistent with the underlying class imbalance.

Remarks.

The pipeline is implemented in Python using PyTorch and HuggingFace Transformers, with scikit-learn for metrics and stratification, pandas and numpy for data handling, and tqdm for instrumentation. Because Governance positives are scarce, the loss uses a high positive weight and the decision threshold is calibrated out-of-sample on validation. Although both tasks share the encoder, their losses are combined using a fixed task weight and that weight is not adapted over training; we also do not ensemble across pooling choices. These choices favor transparency and ease of replication and help explain the predictable precision–recall trade-off observed when moving from the conventional cut-off to the tuned cut-off.