AI Regulation and Entrepreneurship

Junida Mulla
Saïd Business School, University of Oxford

August 2025[1]

PRELIMINARY AND INCOMPLETE

Abstract

AI regulation is being rapidly adopted to manage the trade-offs between societal risks and potential productivity gains. However, its impact on innovation, particularly within the startup ecosystem, remains underexplored. This paper examines how AI regulation affects startups by leveraging the staggered introduction of AI regulation proposals and enactments across 28 U.S. states. The findings show that AI regulation proposals decrease the likelihood of AI startups securing funding by 2.3% annually and increase the probability of acquisition by 1.5%. Using an instrumental variable approach, I confirm that formal regulation enactments further amplify the negative impact on fundraising success. These findings suggest that regulatory pressures may be reshaping startups' access to capital and pushing them toward acquisition as a strategic exit.

Key words: AI regulation, startup financing, exit strategy

[1] I would like to thank Thomas Hellmann, Matthias Qian, Ludovic Phalippou, Alan Morrison, Renee Adams, Jiaqi Zheng, Alex Montag and Erand Mulla for constructive feedback on this draft.

The rapid progress in Artificial Intelligence (AI) has generated both optimism about its potential to boost productivity and significant concerns about its risks to society (Acemoglu, 2024; Gurreiro et al. 2023). As discussions around issues like data privacy, discrimination, and unfair competition intensify, regulatory pressures are mounting. In response, many countries are swiftly adopting new laws and guidelines to address these challenges (Comunale and Manera, 2024). Despite the significant attention these regulatory changes have received, we still have limited evidence into how they might affect innovation and impact startups' ability to adapt and grow.

The debate on AI regulation presents two seemingly contrasting perspectives. On one side, there are concerns that overregulation could stifle innovation (Tartaro et al., 2023), with critics pointing to the immediate effects of the European Union's (EU) regulatory stance on big-tech companies. For instance, Meta delayed the launch of its latest algorithm in the EU due to regulatory uncertainty, while Apple postponed the release of new AI features for similar reasons (Weatherbed, 2024; Campbell, 2024). These responses from industry giants raise concerns about the additional burdens that smaller firms might face under such regulations. In contrast, others argue that AI regulation has minimal negative effect on firms in practice. The rapid pace of AI development often exceeds the federal 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 the firms can potentially bypass or avoid. While both views are theoretically possible, there is a lack of empirical evidence on the actual implementation and effects of AI regulation, leaving the real-world consequences on firms largely unexamined.

In this paper, I explore whether AI regulation influences the fundraising outcomes and exit strategies of AI startups. 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). Fundraising, which provides essential financial resources, is crucial for these startups to grow and develop their technologies. Additionally, Initial Public Offerings (IPOs) and acquisitions offer an economic structure through which these innovations can scale and reach broader markets. Ultimately, the bigger question this paper addresses is how AI regulation impacts innovation through the startup ecosystem, particularly by shaping access to capital and the viability of exit strategies like acquisitions.

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 that local networks play 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 performance data and Bryan Cave Leighton Paisner (BCLP) for information on state-level AI regulations.

I find that the proposal of AI regulation results in a 2% annual decrease in the likelihood of startups securing fundraising, indicating that such regulation may increase perceived risks or compliance costs for investors. However, this effect does not extend to the total amount raised by startups that successfully secure funding, suggesting that while regulation affects access to capital, it does not diminish investor confidence in startups that manage to navigate the regulatory environment. Additionally, I find that regulation proposals increase the probability of a startup being acquired by 1.5%, suggesting that regulation may shape exit strategies, encouraging startups to pursue acquisition as a potential way to bypass regulatory challenges. This shift in exit strategy could reflect startups seeking a more secure route to growth under increased regulatory scrutiny, opting for acquisition by larger firms better equipped to handle the compliance burden.

AI regulation proposals are akin to announcement effects, creating initial uncertainty and fear about potential regulations. Building on this, I also examine the effect of regulation enactments, focusing on states that have already introduced AI regulations, to see how the formal implementation reinforces or escalates the impact beyond the initial proposal phase. As regulations are enacted with a time lag from proposals, difference-in-difference estimates could be biased due to potential anticipation of regulatory enactment. Lobbying pressures might result in weaker regulations being enacted, potentially biasing estimates upward. To account for this, I instrument the likelihood of enactment using the historical partisan composition of the state senate, with more Democratic representation leading to a higher probability of AI regulation enactment. Startups in states with regulation enactments experience a 3% annual drop in the

probability of fundraising compared to startups in states with non-enacted policy proposals. Thus, upon enactment, the effect is 50% stronger than the impact of proposals, indicating that formal enactment amplifies the concerns initially raised during the announcement stage.

Finally, I examine the effect of AI regulation on non-AI startups. Surprisingly, I find that AI regulation negatively affects their probability of securing fundraising, although the impact is less pronounced than for AI-focused firms. Non-AI startups, particularly in sectors like Fintech, may be indirectly affected by regulations around data sharing and privacy. Additionally, AI regulations could limit these startups' ability to adopt AI technologies, reducing their growth prospects and potential for fundraising.

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. Focusing on a specific regulatory approach, Kretschmer et al. (2023) critique the EU's risk-based AI regulation, advocating instead for ex-post liability rules that impose sanctions after harmful outcomes occur. They argue this approach promotes responsible AI use and prevents firms from

bypassing restrictive ex-ante regulations. Ex-ante rules, they note, disproportionately burden developers and favor larger firms that can absorb legal costs, potentially stifling innovation. This is particularly relevant for startups, which may struggle with compliance costs, limiting their ability to compete and innovate under such regulations.

In connection to this theoretical literature, several empirical studies have explored the impact of specific privacy regulations. For instance, 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 instead 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. While these studies provide valuable insights into privacy regulations, the effects of broader AI regulation remain understudied. Additionally, as these studies evaluate a single regulation affecting all firms in a region, challenges arise in ensuring the validity of the parallel trends assumption, particularly with respect to suitability of the control group (Johnson, 2022). This paper addresses these concerns by leveraging the staggered rollout of AI regulations across different U.S. states, comparing AI startups in states with proposed policies to those in

states without. Consistent with these studies, I find that AI regulation negatively affects the probability of funding. Additionally, I show that regulation shifts AI startups' exit strategies, increasing their likelihood of being acquired. Results from Instrumental Variables regressions on the effect of regulatory enactments further reinforce this story.

This paper also relates to the broader empirical research examining the impact of cross-state and cross-country regulations on entrepreneurship. Klapper et al. (2006), for instance, analyze the costs imposed by cross-country entry regulations on new firm creation, particularly in high-entry industries. Similarly, several studies use cross-state variations in tax credits to explore their influence on entrepreneurial activity. Denes et al. (2023) study how tax credits for angel investors affect startup entry and performance, while Fazio et al. (2019) examine the role of R&D tax credits in shaping the quantity and quality of entrepreneurship. More closely related to the current paper, Babina et al. (2024a) exploit the staggered implementation of Open Banking policies across countries to assess their effect on venture capital activity in fintech startups. These studies collectively show how regional regulatory differences shape the entrepreneurial ecosystem by influencing both the rate of new business formation and startup quality. This paper contributes to this literature by focusing specifically on the impact of AI regulation on the financing and exit strategies of startups.

This paper is structured as follows. Section I explains the institutional context behind AI regulation introduction across the U.S. and describes the main data sources used. Section II

explains the empirical methods used to estimate the effect of AI regulation on startup performance outcomes. Section III discusses the main results and section IV 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.[2]

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

At the federal level, the U.S. has not yet passed a comprehensive legislation targeted at regulating AI (BCLP 2024a; White and Case, 2024).[3] 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 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

[3] 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).

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 those on biometric data, facial recognition, 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 legislation. 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.[4]

Insert Figure I about here

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 regulations, with New York leading the count at 20 proposals[5]. 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. Of the 79 proposed policies, only 19 have been enacted across 16 states. On average, it takes 277 days for regulations to move from proposal to enactment.

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

[4] 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).

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

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.[6] 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 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.

[6] 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.

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. Empirical strategy

This section outlines the empirical strategy used to examine the relationship between AI regulation and startup performance. In Section II.A, I employ a staggered Difference-in-Differences (Diff-in-Diff) approach to assess the role of regulation proposals across states that adopt AI policies at different times. Section II.B describes the instrumental variable approach to examine the effect of policy enactments on startup performance.

A. Policy Proposals

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. 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.[7] 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). [8] I therefore adopt an Instrumental Variable (IV) approach to study the effect of enacting AI regulation.

[7] 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).

[8] 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.

Using a methodology similar in spirit to Acemoglu et al. (2023), 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.[9] Since Democrats have traditionally been more 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

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

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:

$$y_{ist} = \beta^{en} Enactment_s \times Post_{st} + ElectionYr_{st} + Age^{en}_{ist} + \delta^{en} Z_{st-1} + \alpha^{en}_{is} + \alpha^{en}_{t} + \xi_{ist} \quad (3)$$

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:

$$Enactment_s \times Post_{st} = DemShare_s \times Post_{st} + \alpha^{en}_{is} + \alpha^{en}_{t} + Age^{en}_{ist} + \delta^{en} Z_{st-1} + \xi_{ist} \quad (4)$$

III. Results

This section presents the main findings. Section III.A explores the relationship between regulation proposals and startup fundraising outcomes. Section III.B discusses the

robustness of these results. Section III.C extends the analysis to exit strategies, while Section III.D examines the effect of policy enactments. Finally, Section III.E shifts the focus to non-AI startups.

A. 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 5%) on the probability of fundraising success amounting to a 2.3% annual decrease. 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

Several mechanisms could explain the negative effect of AI regulation on the probability of fundraising. AI regulations may restrict a startup's ability to collect and sell data, increase compliance costs, or impose direct financial penalties for violations (for a discussion of these channels see e.g., Canaya et al. 2022; Tartaro et al. 2023). I intend to explore these further in the upcoming versions of the paper, by exploiting the variation in the type of regulations different states propose.

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.[10] 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 the descriptions of both the companies and the regulations, it becomes clear that the similarity scores accurately reflect the relevance of the regulations, correctly identifying which regulations are most and least 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.

B. Robustness Checks

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

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.

C. 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 Difference-in-Differences 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.

D. Effect of Policy Enactment

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.

Table V shows OLS and 2SLS estimates, following the methodology described in Section II.B. Regressions (1), (3), (5), and (7) display OLS estimates from the uninstrumented equation (3), while regressions (2), (4), (6), and (8) 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 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% annual decrease in the probability of

fundraising, which is larger in magnitude than the 2% decrease associated with regulation proposals.

Enactments also positively affect the probability of acquisition, in line with the results in Table IV. Interestingly, the direction of bias in column (7) is opposite to that in column (1), with the OLS coefficient being smaller than the 2SLS estimate. While this might seem counterintuitive at first, it makes economic sense: if lobbying leads to the enactment of less harmful regulations, startups face lower compliance costs and therefore have less need to be acquired.

E. Effect on the rest of the start-ups

Having established the detrimental role of AI regulation on the probability of fundraising for AI startups, I next explore whether similar patterns hold for the remaining non-AI startups. As described in Section I, AI startups are identified by a set of specific keywords. Any startup which does not contain any of the keywords in the startup description or industry classification is defined as a non-AI startup for the purposes of this section. In Table VI, I compare the results between AI and non-AI startups. Columns (1) and (2) display TWFE estimators, whereas columns (3) and (4) show Sun and Abraham (2021) estimates for the role of policy proposals on the probability of fundraising. Columns (5)-(8) display the effect of policy enactments, with OLS coefficients displayed in columns (5) and (6) and 2SLS coefficients displayed in columns (7) to (8).

The negative effect of AI regulation on the probability of fundraising persists in non-AI startups, though the magnitude is roughly half that of AI startups. The effect is very stubborn, and does not go away when incorporating additional controls, such as VC activity or broader measures of financial climate within states.

There are several potential explanations for the negative effect observed in non-AI startups. First, startups outside the AI sector, such as Fintech, may be indirectly affected by AI regulations, particularly where data sharing, privacy, or similar activities overlap with their operations. The general regulatory tightening around technology could create investor uncertainty, limiting fundraising opportunities across the broader ecosystem. Finally, the onset of AI regulation may prevent non-AI startups from investing in AI, reducing their growth potential. This is in line with randomized survey evidence suggesting that providing managers with information regarding AI regulation lowers their willingness to adopt AI (Cuellar et al., 2024).

IV. 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 and exit strategies such as acquisitions.

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% annually, reflecting heightened perceived risks or increased compliance costs for investors. 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%, suggesting that acquisitions may become a preferred strategy to navigate regulatory hurdles. These trends persist even after the enactment of regulations. Interestingly, similar effects are observed for non-AI startups, indicating that AI regulation may have wider implications, potentially discouraging broader AI adoption.

References

Acemoglu, D. (2021). *Harms of AI* (Working Paper No. 29247). National Bureau of Economic Research.

Acemoglu, D., He, A., & Le Maire, D. (2022). *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.*

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).

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.

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.

Bureau of Economic Analysis. 2024. *BEA Data*. Available at: https://www.bea.gov/data (Accessed 18 August 2024).

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).

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-bbe6-14fb6e85b4cc#:~: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).
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)
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.
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.
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.
Eldar, O., & Grennan, J. (2024). Common venture capital investors and startup growth. *The Review of Financial Studies,* 37(2), 549-590.
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.
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.*
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.
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).
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 Economic*s, 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.
Kalmenovitz, J. (2023). Regulatory intensity and firm-specific exposure. *The review of financial studies*, 36(8), 3311-3347.
Klapper, L., Laeven, L., & Rajan, R. (2006). Entry regulation as a barrier to entrepreneurship. *Journal of Financial Economics*, 82(3), 591-629.
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.
Lee, I. H. (2022). Startups, relocation, and firm performance: a transaction cost economics perspective. *Small Business Economics,* 58(1), 205-224.
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).
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)
Open AI. 2022. *New and improved embedding model.* Available at: https://openai.com/index/new-and-improved-embedding-model/ (Accessed 5 October 2024).
Samuel, J., Khanna, T., & Sundar, S. (2024). F*ear 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.

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.

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).

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).

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.*

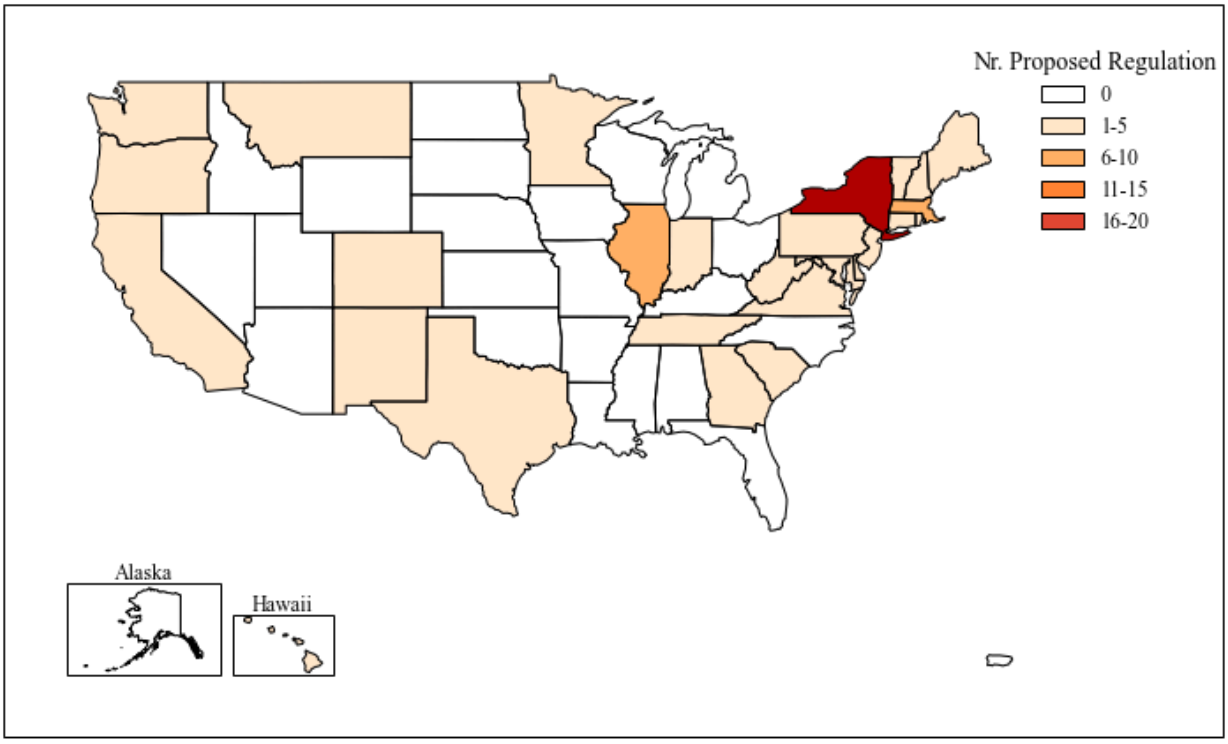
White & Case. (2024). *AI Watch: Global regulatory tracker - United States*. Available at: https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-united-states#:~:text=Currently%2C%20there%20is%20no%20comprehensive,prohibit%20or%20restrict%20their%20use (Accessed 2 October 2024)

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

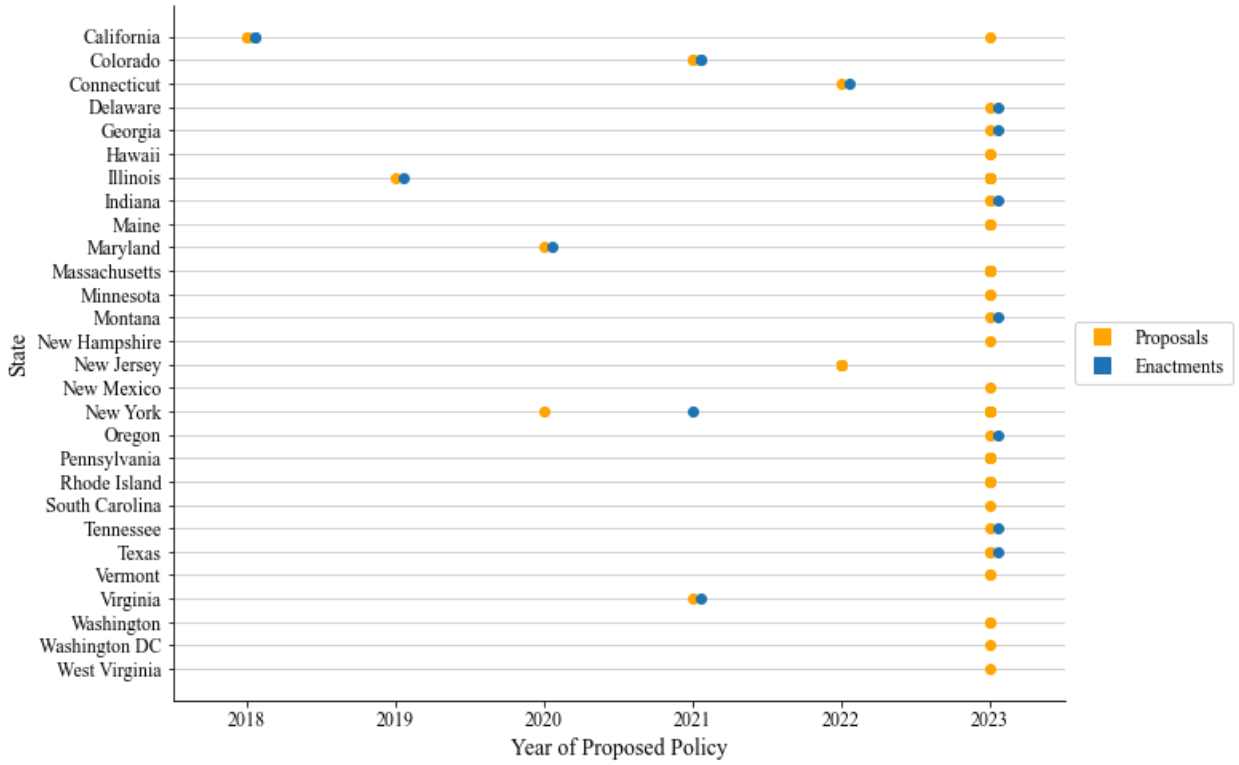
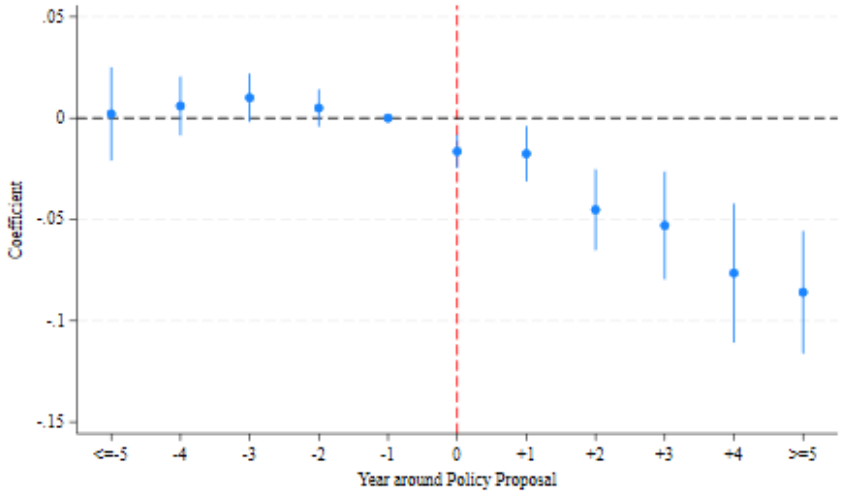


Figure II: Dynamic Effect of AI regulation proposals

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)

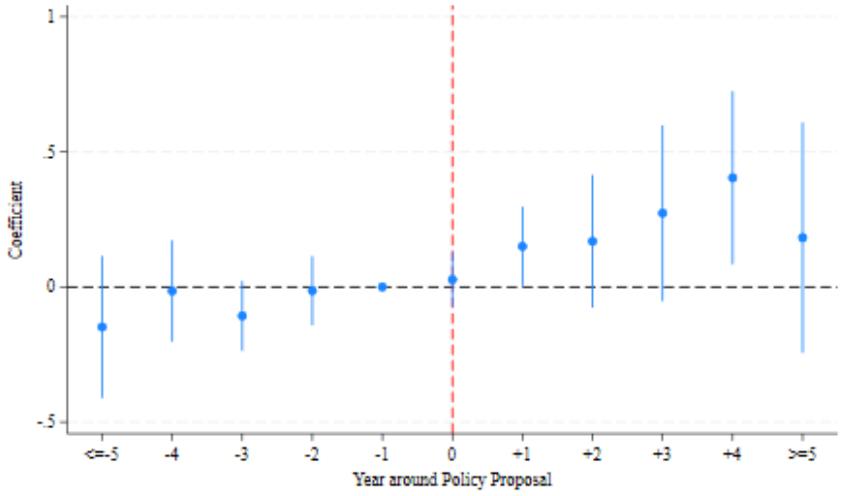


Table I: Variable Definitions

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-emebedding-ada-002 model. Source:
Enacted	A dummy variable equal to 1 if the state in which the company is located enacted its first AI regulation in a particular year. Dummy remains 1 for all years subsequent to initial enactment. Source: BCLP.
DemShare	The 6-year average of the proportion of Democratic state senate members, calculated over the 6-year period leading up to the first AI regulation proposal in states that introduced such regulations. Source: Ballotpedia
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 logarith 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.
Election Year	A set of dummy variables indicating each calendar year in which a state held senate elections. Source: Ballotpedia.

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. The number of observations for DemShare decreases due to missing data.

	Obs.	Mean	St. Dev.	p5	p50	p95
			Company-year level			
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
			State-year level			
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
DemShare	598	0.4388	0.2059	0.1515	0.4082	0.8649

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***		0.0775*	
	(0.0056)		(0.0436)	
Exposure		-0.0302***		0.0919
		(0.0077)		(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		0.0145***	
	(0.0002)		(0.0014)	
Exposure		0.0005		0.0195***
		(0.0003)		(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: AI regulation enactments and startup performance

This table presents OLS and 2SLS regression estimates for the effect of AI regulation enactments on startup outcomes, for startups in states with proposed regulation. Regressions (1), (3), (5), and (7) display estimates from the uninstrumented equation (3), while regressions (2), (4), (6), and (8) show 2SLS estimates, where enactments are instrumented by the 6-year average proportion of Democratic senate members, as described in Section II.B. 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. Columns (5) and (6) use a binary indicator for whether a startup had an IPO, while columns (7) and (8) measure whether a startup was acquired. 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)		DV: IPO		DV: Acquired	
	OLS	IV Estimate	OLS	IV Estimate	OLS	IV Estimate	OLS	IV Estimate
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Enacted	-0.0138*	-0.0316***	0.0938	0.1271**	0.0002	0.0001	0.0095***	0.0127***
	(0.0076)	(0.0104)	(0.0628)	(0.0591)	(0.0003)	(0.0005)	(0.0030)	(0.0031)
Time-varying state controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Election Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	84,928	83,695	8,331	8,260	84,691	83,459	80,686	79,476
Adjusted R^2 (Centered for IV)	0.2439	0.0014	0.7095	0.0026	0.0180	0.0002	0.0615	0.0017
	First-Stage for IV Regressions							
DemShare		1.2418***		1.3764***		1.2428***		1.2527***
		(0.2024)		(0.1315)		(0.2018)		(0.1964)
F-Stat		37.6297		109.5393		37.9389		40.6686

Table VI: AI regulation and non-AI start-ups

This table examines the impact of AI regulation on the fundraising outcomes of non-AI startups. Non-AI startups are defined as those whose descriptions do not contain any of the AI keywords listed in Section I.B. Panel A presents TWFE DiD and SA DiD estimates for the effect of AI regulation proposals, based on equation (1). Panel B shows OLS and 2SLS estimates for the effect of AI regulation enactments, where enactments are instrumented using the 6-year average proportion of Democratic senate members prior to the first regulation proposal, as outlined in Section II.B. In Panel B, the sample is restricted to startups in states with regulation proposals. For all regressions (1)-(8), the dependent variable is a binary indicator of whether a startup raised funding in a given year. Time-varying state controls include the lagged (logged) level and growth rate of real GDP, along with lagged (logged) disposable income per capita. Standard errors are clustered at the state level, and significance is indicated by ***, **, and * for 99%, 95%, and 90% confidence levels, respectively.

Panel A: Effect of Policy Proposal on Probability of fundraising for AI start-ups vs Non-AI startups				
	TWFE DiD Estimator		SA DiD Estimator	
	(1)	(2)	(3)	(4)
	AI Startups	Non-AI Startups	AI Startups	Non-AI Startups
Proposed	-0.0231***	-0.0162***	-0.0520***	-0.0243***
	(0.0056)	(0.0029)	(0.0068)	(0.0021)
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	2,008,461	99,729	2,008,461
Adjusted R^2	0.2479	0.3083	0.2492	0.2168
Panel B: Effect of Policy Enactment on Probability of fundraising for AI start-ups vs Non-AI startups				
	OLS		IV Estimator	
	(5)	(6)	(7)	(8)
	AI Startups	Non-AI Startups	AI Startups	Non-AI Startups
Enacted	-0.0138*	-0.0080**	-0.0316***	-0.0182***
	(0.0076)	(0.0031)	(0.0104)	(0.0040)
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	1,511,780	83,695	1,489,225
R^2 (Centered for IV)	0.2439	0.2214	0.0014	0.0009
			First-Stage for IV regressions	
DemShare			1.2418***	1.3183***
			(0.2024)	(0.1495)
F-Stat			37.6297	77.8015

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 than 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: Additional Tables and Figures

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 Callaway and Sant'Anna (2021) (CS DiD) estimator, whereas columns (3) and (4) report results using the Sun & Abraham (2021) (SA 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)	(2)	(3)	(4)
	Raised Eq	Amount Raised (log)	Raised Eq	Amount Raised (log)
Proposed	-0.0520***	0.0653	-0.0579***	0.3724
	(0.0068)	(0.0913)	(0.0215)	(0.8508)
Time-varying state contro	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***	0.1039
	(0.0064)	(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***	0.0838
	(0.0051)	(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	0.0476
	(0.0087)	(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	0.1375*
	(0.0086)	(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	0.0673
	(0.0129)	(0.1525)
Controls	Yes	Yes
Observations	38,484	3,055
Adjusted R^2	0.2556	0.6829