

**HOW WOULD AI REGULATION CHANGE
FIRMS' BEHAVIOR?
EVIDENCE FROM THOUSANDS OF
MANAGERS**

Yong Suk Lee
Stanford University

Michael Webb
Stanford University

Benjamin Larsen
Copenhagen Business
School

**Mariano-Florentino
Cuéllar**
California Supreme Court &
Stanford University

November, 2019
Working Paper No. 19-031

How Would AI Regulation Change Firms' Behavior? Evidence from Thousands of Managers

Yong Suk Lee, Stanford University

Benjamin Larsen, Copenhagen Business School

Michael Webb, Stanford University

Mariano-Florentino Cuéllar, California Supreme Court and Stanford University

November 2019

ABSTRACT

We examine the impacts of different proposed AI regulations on managers' intentions to adopt AI technologies and on their AI-related business strategies. We conduct a randomized online survey experiment on more than a thousand managers in the U.S. We randomly present managers with different proposed AI regulations, and ask them to make decisions about AI adoption, budget allocation, hiring, and other issues. We have four main findings: (1) information about AI regulation generally reduces the rate of adoption of AI technologies. Nonetheless, industry- and agency-specific AI regulation has a smaller impact than general AI regulation. (2) Information about regulation induces firms to think. That is, firms spend more on developing AI strategy and hire more managers. This is at the cost of hiring other workers and training current employees. (3) The impact of information about AI regulation on innovation differs by industry and firm size. AI regulation increases intent to file patents in the healthcare and pharmaceutical sectors, but reduces it in the retail sector. Moreover, AI regulation information reduces AI adoption in small firms and is more likely to reduce their innovative activity. (4) Information about AI regulation increases firms' perceptions of the importance of safety and transparency issues related to AI.

Keywords: artificial intelligence, regulation, adoption, innovation, strategy, survey experiment

JEL Codes: K24, L21, L51, O33, O38

1. Introduction

Artificial intelligence (AI) technologies have become increasingly widespread over the last decade. In particular, the fields of image recognition, speech recognition, data analytics, and machine translation have advanced rapidly, spurred by important breakthroughs in deep neural networks (Varian, 2018). But as the use of artificial intelligence has become more common and the performance of AI systems has improved, policymakers, scholars, and advocates have also raised concerns. Issues of algorithmic bias, data privacy, and transparency have gained increasing attention as a result, raising renewed calls for policy efforts to address the consequences of technological change (Frank et al., 2019). As AI continues to improve and diffuse, it will likely also have important long-term consequences for jobs, inequality, organizations, and competition. These developments may spur interest in regulation as a potentially important means for addressing the risks and possibilities of AI.

Yet very little is known about how different kinds of AI-related regulation — or even the prospect of regulation — might affect firm behavior. AI is already being implicitly regulated through common law doctrines such as tort and contract law, as well as statutory and regulatory obligations on organizations, such as emerging standards governing autonomous vehicles (Cuéllar 2019). As AI technologies are diffusing rapidly and have wide-ranging social and economic consequences, policymakers as well as federal and state agencies are also contemplating new ways of regulating AI. These include broad proposals of general AI regulation such as the Algorithmic Accountability Act, which was introduced in The House of Representatives on April 10, 2019. State regulations include the California Consumer Privacy Act, which goes into effect from January 2020. Domain-specific regulations are also currently being developed by federal regulators such as the Food and Drug Administration (FDA), the

National Highway Traffic and Safety Administration (NHTSA), and the Federal Trade Commission (FTC).

In this paper, we examine the impact of these actual and potential AI regulations on business managers. In particular, we assess how likely managers are to adopt AI technologies and alter their AI-related business strategies as they are asked to reflect on the regulation of AI. We conduct a randomized online survey experiment where the treatment group is informed of the core features of different regulatory treatments. Specifically, we randomly expose managers to one of the following treatments: (1) a general AI regulation treatment that invokes the Algorithmic Accountability Act, (2) industry-specific regulation treatments that invoke the relevant agencies, i.e., the FDA (for healthcare, pharmaceutical, and biotech), NHTSA (for automobile, transportation, and distribution), and the FTC (for retail and wholesale), (3) a treatment that reminds managers that AI adoption in businesses are subject to existing common law and statutory requirements such as tort law, labor law, and civil rights law, and (4) a data privacy regulation treatment that invokes the California Consumer Privacy Act. Specifically, we study how these varying regulations affect managers' decision-making, and how managers revise their business strategies when faced with new regulation.

Our results indicate that exposure to information about regulation decreases managers' reported intent to adopt AI technologies in the firm's business processes. We find that exposure to information about general AI regulation, such as the Algorithmic Accountability Act, reduces the reported number of business processes in which managers are reportedly willing to use AI by about 16%. We also find that exposure to information about AI regulation significantly increases expenditure on developing AI strategy. This impact is strongest for the general AI regulation treatment, which increases allocation to AI strategy purposes by 3 percentage points. The

increase in budget for developing AI business strategy is primarily offset by a decrease in the budget for training current employees on how to code and use AI technology, and purchasing AI packages from external vendors. In other words, making the prospect of AI regulation more salient seems to force firms to “think,” inducing managers to report greater willingness to expend more on strategizing, but at the cost of developing internal human capital.

Exposure to information about AI regulation also increases how importantly managers consider various ethical issues when adopting AI in their business. Each regulation treatment increases the importance managers put on safety and accident concerns related to AI technologies, and the existing AI regulation and data privacy regulation treatment significantly increase manager perceptions of the importance of privacy and data security. The agency-specific regulation also increases manager perceptions of the importance of bias and discrimination, and transparency and explainability.

We find significant heterogeneity in the impact of AI regulation information by industry and firm size. Regulation decreases AI adoption in the healthcare and retail sectors but not the transportation sector. Moreover, it is primarily in the transportation sector that AI regulation results in higher budget allocation to developing AI strategy. In terms of innovation activities, we find that AI regulation increases firms’ intent to file patents in the healthcare sector but decreases it in the retail and wholesale sector. This is likely due to patents being a vital part of the healthcare industry (i.e. drug discovery), while the core business in retail is far less dependent on patents as a primary strategy for operation. The negative impact of AI regulation information on AI adoption is more significant for small firms, which we define as those with revenue less than \$10 million. Also, these small firms are the ones that increase their budget allocation to AI strategy and hire more managers in response to new regulations. However, large firms respond to

the existing AI regulation treatment, which invokes the relevance of tort law and civil rights protections. Managers of large firms exposed to this treatment increase their awareness of ethical issues, increase the budget share for developing AI strategy, and plan to hire more managers. These results highlight the potential trade-offs between regulation and the diffusion and innovation of AI technologies in firms, and provide important implications for regulators and policymakers.

To the best of our knowledge, our paper is the first to examine the potential impact of AI regulation on reported AI adoption and innovation. Our findings are especially relevant to the literature that examines the effects of technology-related regulations, especially privacy regulation. In this line of research, Goldfarb and Tucker (2012) have found that in data-driven industries, privacy regulation impacts the rate and direction of innovation. Too little privacy protection means that consumers may be reluctant to participate in market transactions where their data are vulnerable. Too much privacy regulation means that firms cannot use data to innovate. The evidence generally indicates that most attempts at government-mandated privacy regulation lead to slower technology adoption and less innovation (Goldfarb and Tucker, 2011; Miller and Tucker, 2011; Kim and Wagman, 2015).

Another related aspect is the liability risk of AI, such as that arising from algorithmic bias. While researchers have found evidence suggesting that AI predictions are less biased than human predictions, they may still perpetuate biases present in the data used to train them (Hoffman, Kahn, and Li, 2018; Fuster et al. 2018). Furthermore, it is easier to audit AI-based decisions than human decisions, which opens up the possibility of an increase in liability claims. Firms face liability risk even in situations when the bias is unintended (Agrawal, Gans, and Goldfarb, 2018b). Such risk could serve as a deterrent to the adoption of AI technologies.

We also contribute to the literature on the diffusion of new technologies. Machine learning technologies have not yet been widely diffused, which means that the full effects of AI technologies will not be realized until waves of complementary innovations are developed and implemented (Brynjolfsson, Rock, and Syverson 2017). In particular, business process redesign, co-invention of new products and business models, and investments in human capital are likely all needed before the economy can experience significant AI-driven productivity gains (Brynjolfsson, Rock, & Syverson, 2018). These findings suggest that policy should be dealing not only with the consequences of AI, but also with how to support its effective incorporation and, where appropriate, ongoing diffusion.

Finally, there is a rapidly growing literature on the potential labor market consequences of automation from AI and robotics (e.g., Aghion, Jones, and Jones 2017; Brynjolfsson et al. 2018; Acemoglu and Restrepo 2019a, 2019b; Lee and Chung 2019; Webb 2019; Dixen et al. 2019). This literature finds that automation may lead to declines in employment and wages, at least in the short run, but may increase employment in the long run. The literature also suggests that the effects of automation may likely be different for different occupations. Our finding that AI regulation may result in a reduction in AI-related training within firms suggests that AI regulations may have direct impacts on labor markets, as well as on AI diffusion and the rate of innovation.

The outline of the paper is as follows. In the next section, we provide background on the current state and potential directions of AI regulation. Section 3 discusses the empirical strategy, and Section 4 the data and sample. In section 5 we report our main results, followed by an exploration of heterogeneous impacts in Section 6. In section 7 we offer some concluding comments.

2. AI Regulation

AI describes a broad set of computing techniques and associated technologies with widespread applications in a variety of workplace, commercial, and governmental settings. This makes it hard to generalize about the most desirable rules to govern real-world applications of AI and human interaction with AI. The incorporation of AI in autonomous vehicles may for instance call for a variety of standards governing road safety, inter-vehicle communication, ethical dilemmas and cybersecurity, while use of AI in healthcare or retail may call for altogether different standards focusing on matters such as privacy and disclosure. Other domains of use, such as the utilization of AI in hiring decisions, in the judicial system, in aviation, and so on, may call for rules and regulations to promote accountable, unbiased and safe application. For many observers, a call to heighten regulation of AI is almost certainly spurred by an increase in the use of AI technologies, combined with a perceived lack of control and oversight of existing AI practices. Public perceptions of the relationship between individual economic well-being and the generation of data are slowly changing, however, as evidenced in proposals for policies such as a “data dividend,” where companies would have to pay for consumers’ data (Cuéllar and Huq 2019).

To better understand how business managers respond to concerns about AI regulation, in this paper we cover six existing and tentative approaches to AI regulation, moving from existing laws and statutes to the proposed Algorithmic Accountability Act, and the incoming California Consumer Privacy Act. We also cover three domain-specific approaches across healthcare, automotive, and retail. Our central goal is to understand how different regulatory approaches, current and intended, will have an impact on businesses’ rate of AI adoption and innovation across varying industries.

Even as policymakers consider how to adapt to the growing ubiquity of AI, the regulatory landscape is already quite broad. Because existing laws apply broadly to individuals and organizations regardless of whether they rely on AI or not, in the United States the use of AI is already implicitly governed by a variety of common law doctrines and statutory provisions, such as tort law, contract law, and employment discrimination law (Cuéllar 2019). This implies that judges' rulings on common law-type claims already plays an important role in how society governs AI. While common law builds on precedence, federal agencies also engage in important governance and regulatory tasks that may affect AI across a variety of sectors of the economy (Barfield and Pagollo 2018). Federal autonomous vehicle legislation, for instance, carves out a robust domain for states to make common law decisions about autonomous vehicles through the court system. Through tort, property, contract, and related legal domains, society shapes how people utilize AI, while gradually defining what it means to misuse AI technologies (Cuéllar 2019). Existing law (e.g., tort law) may, for instance, require that a company avoid any negligent use of AI to make decisions or provide information that could result in harm to the public. Likewise, current employment, labor, and civil rights laws imply that a company using AI to make hiring or termination decisions could face liability for its decisions involving human resources.

Policymakers and the public nonetheless often consider new legal and regulatory approaches when faced with potentially transformative technologies because these technologies may pose challenges for, and ultimately fail to fit the purpose of, existing laws and regulations (Barfield and Pagollo 2018). The Algorithmic Accountability Act is one such proposal. Co-sponsored by several federal legislators, the Act would regulate large firms with gross annual receipts of \$50 million or more over the last three consecutive years, or which possess or control

personal information on more than 1 million consumers (Congress 2019). Other possibilities include the establishment of an Artificial Intelligence Regulation Agency that is independent of federal regulators (Weaver 2018). Another approach is suggested by Clark and Hadfield (2019), in which regulation is outsourced to regulatory markets, while oversight is handled by private regulators in concert with government and policymakers. Until now, so-called “soft” law governance, such as The Partnership on AI, as well as IEEE standards addressing governance and ethical aspects of AI, continue to play a role in setting the default for how AI is governed (Wallach and Marchant 2018).

While statutes imposing new regulatory requirements such as the Algorithmic Accountability Act are still under debate, regulation of data privacy is already being implemented. The state of California recently introduced the California Consumer Privacy Act (CCPA), which goes into effect in January 2020. The CCPA will affect all businesses buying, selling, or otherwise trading the “personal information” of California residents, including companies using online-generated data from residents across their products. The CCPA thus adds another layer of oversight to the area of data handling and privacy, on which many AI applications are contingent.

Although the common law, existing statutes, and forthcoming privacy regulations already govern many terms of usage related to AI application and data handling, domain-specific regulators are also devising their own approaches to regulate AI. In this study, we have chosen to focus on the current regulatory approaches to healthcare, automotive, and retail, and so focus on the current initiatives applied by the Food and Drug Administration (FDA), the National Highway Traffic and Safety Administration (NHTSA), and the Federal Trade Commission (FTC).

In the spring of 2019, the Food and Drug Administration (FDA) released a ‘Proposed Regulatory Framework for Modifications to AI/Machine Learning Based Software as a Medical Device.’ The FDA’s approach to regulate AI aims to examine and pre-approve the underlying performance of a firm’s AI products before they are marketed, as well as post-approving any subsequent algorithmic modifications. The proposed regulatory framework takes into consideration a total product lifecycle-approach in which AI technologies and products will remain open to real-world learning and adaptation through continuous algorithmic updating, while ensuring that standards for safety and efficiency are met.

The National Highway Traffic and Safety Administration (NHTSA) regulates the autonomous vehicle and logistics industry, and has emphasized the importance of removing unnecessary barriers to innovation. NHTSA has for instance specified that its current safety standards for Level 4 and 5 automated vehicles (so called, “fully autonomous” vehicles) constitute an unintended regulatory barrier to innovation, while existing regulations and vehicle safety standards will remain in effect until a revised framework for automated driving systems is established. The approach taken by NHTSA exemplifies a light-touch approach to AI regulation, which provides ample space for innovation in autonomous vehicle technologies.

The Federal Trade Commission (FTC) is the primary agency responsible for regulating e-commerce activity, which includes online advertising, consumer privacy, and commercial emails. Since AI is being heavily used in e-commerce and online marketing, the FTC has engaged in a series of fourteen ‘Hearings on Competition and Consumer Protection in the 21st Century,’ to safeguard consumers from unfair and deceptive practices. Some of these hearings in late 2018 focused on ‘Algorithms, AI and Predictive Analytics,’ ‘Privacy, Big Data and Competition,’ and ‘Data Security’. As the retail sector has been especially fast at deploying and

monetizing a range of AI technologies on online and e-commerce platforms, revamped oversight by the FTC is likely to require firms operating in the space to assess and disclose the impact of their AI systems on various issues. The hearings concluded in June 2019, and it remains to be seen what kind of initiatives may emerge from them.

We have seen, in short, that AI regulation is emerging and is likely to materialize more intensely across several directions simultaneously: from existing laws, new general regulations, and evolving domain-specific regulations. The main goal of regulators is to ensure opportunity in the application and innovation of AI-based tools, products, and services while limiting negative externalities in the areas of competition, privacy, safety, and accountability. It remains little known, however, how the proposed Algorithmic Accountability Act, the incoming CCPA, as well as the regulatory approaches taken by the FDA, NHTSA, and the FTC, will affect the rate of AI adoption and innovation across different firms and industries.

3. The Online Survey Experiment

We conduct a randomized online survey experiment to study the effects of different regulatory treatments on three broad industries: healthcare/pharmaceutical/bio-tech (henceforth, healthcare), automotive/transportation/distribution (henceforth, automotive), and retail and wholesale. Specifically, we randomly expose managers in each of these industries to one of the following treatments: a general AI regulation treatment that invokes the proposed Algorithmic Accountability Act (T1); industry-specific regulation treatments that invoke the relevant agencies, i.e., the FDA (for healthcare, pharmaceutical, and biotech), NHTSA (for automotive, transportation, and distribution), and the FTC (for retail and wholesale) (T2); a treatment that reminds managers that AI adoption in businesses are subject to existing common law and

statutory requirements such as tort law, labor law, and civil rights law (T3); and a data privacy regulation treatment based on the incoming (January 2020) California Consumer Privacy Act (T4).

In T2, managers are exposed to one of the three industry-specific treatments based on the nature of their firm's business. The three treatments in T2 correspond to the substantive focus, legal authority, and current approach taken by regulators with authority over particular kinds of business activities, i.e., the FDA, NHTSA, and FTC. The other treatments (T1, T3, and T4) are industry-agnostic and all managers in the treatment group receive the same treatment regardless of the nature of their firm's business. Figure 1 summarizes the structure of the online experiment. Other than for the agency-specific AI regulation treatment, managers in different industries are exposed to the same general AI regulation, existing AI-related regulation, and data privacy regulation statements.

To begin, we present both the treatment and the control groups with an introductory paragraph that contains details about the current and forecasted adoption of AI technologies¹:

“Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.”

For our control group, we seek to balance the preceding paragraph to make it represent some of the same concerns that our treatment group is subjected to, although without specifically mentioning regulation or any form of regulatory compliance.

“While the potential for AI is vast, most organizations still have a long way to go in developing the core practices that enable them to realize the potential value of AI at scale. Business executives and managers will need to think about how to incorporate AI into their business strategy, as well as the transparency and

¹ The contents of the introductory paragraph are based on a McKinsey Global Survey of AI adoption (McKinsey 2018).

“explainability” of AI algorithms, biases in data, and concerns about safety and privacy.”

For the treatment groups, we rephrase the second paragraph (depending on the treatment group) to contain details about the following laws or agencies that could affect the use of AI:

1. The Algorithmic Accountability Act (T1=General Regulation)
2. Food and Drug Administration (T2a=Healthcare Regulation)
National Highway Traffic and Safety Administration (T2b=Automotive Regulation)
Federal Trade Council (T2c=Retail Regulation)
3. Existing Laws (T3=Common Law and Existing Statutes)
4. The California Consumer Privacy Act (T4=Data Privacy Regulation)

For T1 (General Regulation) we stress that the Algorithmic Accountability Act requires firms to disclose their usage of AI systems, including their development process or contractor of origin, AI system design, model training, as well as data gathered and in use. We also note that the Act requires firms to disclose to a government agency the impact of their AI systems on safety, accuracy, fairness, bias, discrimination, and privacy.

For T2a (Healthcare Regulation), we note that the FDA aims to examine and pre-approve, consistent with its legal authority, the underlying performance of a firm’s AI products before they are marketed, and post-approve any algorithmic modifications. We note that the FDA will assess a firm’s ability to manage risks associated with issues such as transparency and explainability (e.g., diagnosis recommendation algorithms), and security (e.g., use and protection of patient private information) of the AI/Machine Learning based software.

For T2b (Automotive Regulation) we specify that NHTSA emphasizes the importance of removing unnecessary barriers while issuing voluntary guidance rather than regulations that

could stifle innovation. We further note that NHTSA has specified that its current safety standards constitute an unintended regulatory barrier to innovation of autonomous driving vehicles, but that existing regulations and vehicle safety standards remain in effect until a revised framework for automated driving systems is established.

For T2c (Retail Regulation) we convey that the FTC has engaged in hearings to safeguard consumers from unfair and deceptive practices surrounding potential issues across algorithmic discrimination and bias (e.g. in online adds / micro-targeting of consumer groups), transparency (e.g. product recommendation engines) and security (e.g. use and protection of consumers private information). We note that revamped oversight by the FTC will likely require retailers deploying AI technologies to assess and disclose the impact of their AI systems across those issues.

For T3 (Common Law and Existing Statutes) we stress that firms using AI technology in the United States are already subject to some common law and statutory requirements relevant to AI. We note that existing laws (e.g., tort law) may require that a company avoid any negligent use of AI to make decisions or provide information that could result in harm to the public. We also remark that current employment, labor, and civil rights laws create the risk that a company using AI to make hiring or termination decisions could face liability for its decisions involving human resources.

For T4 (Data Privacy Regulation), we stress that the California Consumer Privacy Act of 2018 (CCPA) will affect all businesses buying, selling or otherwise trading the “personal information” of California residents - including companies using online-generated data from residents across their products. We note that in order to stay compliant with the regulation, firms must disclose how they use and store personal data, and how they conform with data privacy

rules. Finally, we add that other states are expected to enact similar data privacy regulations in the near future.

For most treatments, except T2b (Automotive Regulation) and T3 (Common Law and Existing Statutes), we set the year on which new regulation takes effect at 2020, in order to minimize variation based on different manager assumptions about the effective date of new regulations. We omit this mention of time when discussing existing laws, while NHTSA's regulatory approach of removing unnecessary barriers to regulation does not warrant a future date of action or implementation. The full texts of the treatments can be found in Appendix Table 1.

Following the treatment/control scenario, participants are asked five sets of questions related to managers' inclination towards 1) adoption of AI technologies; 2) budget allocation; 3) AI-related innovation; 4) ethical issues; and (5) labor. The adoption of AI technologies (i.e. machine learning, computer vision, and natural language processing) is measured as the number of business processes, going from one to ten processes, with a higher number of processes implementing AI signaling a higher degree of AI implementation and usage.

We then ask managers how they would allocate budgets across six expense categories. By forcing the allocation to add to 100 percent, we are able to examine the trade-offs managers choose in response to the perceived impact of AI regulation. We measure budget allocation by having managers fill out six different categories with costs related to: 1) R&D related to creating new AI products or processes; 2) hiring managers, technicians, and programmers, excluding R&D workers, to operate and maintain AI systems; 3) AI training for current employees; 4) purchasing AI packages from external vendors; 5) computers and data centers, including

purchasing or gathering data; and 6) developing AI strategy that is compatible with the company's overall business strategy.

Innovation is addressed by asking managers how likely they are to adjust AI-related innovation activities at their workplace in the coming year across three categories. These are: 1) co-operation on AI-related R&D activities with other institutions such as universities, research institutes, and other businesses; 2) filing of AI-related patents; and 3) introduction of an AI-related good, service, or production/delivery method that is new or improved. We measure managerial adjustments on a standard Likert scale.

Managerial values and ethical issues are assessed by asking the degree of importance that managers attach to: 1) layoffs or labor related issues due to AI adoption; 2) racial and gender bias/discrimination from AI algorithms; 3) safety and accidents related to AI technologies; 4) privacy and data security issues related to AI adoption; and 5) transparency and explainability of AI algorithms. We measure managerial values on a standard Likert scale ranging from not important to very important. In a following question, we ask managers whom they consider to be primarily responsible for AI-related ethical issues in their business: 1) managers; 2) engineers; 3) AI package vendors; 4) the government, i.e., regulatory agencies; 5) the courts; and 6) other.

Finally, we look at labor by asking managers and executives how they would adjust the total number of employees at their workplace across: 1) managers; 2) technical workers, including R&D workers; 3) office workers; 4) service workers; 5) sales workers; and 6) production workers. We specify that we are only interested in changes that would occur because of AI adoption at the workplace²

² The survey can be accessed online at https://web.stanford.edu/~yongslee/AIReg_FDA.pdf, which features the FDA treatment for the healthcare sector. The survey questions for the automotive and retail sectors are the same as above, except for the industry-specific regulation treatment texts, which are presented in Appendix Table 1.

4. Sample and Data

We recruit managers in the US using SurveyMonkey Audience. We focus on managers in businesses of at least 50 employees, since they are likely to be well-aware of the types of technologies being used at their businesses and be involved in the decisions surrounding adoption. The managers we recruited include owners and partners of businesses, C-level executives, and senior and middle managers in the three broad industries discussed above. We launched the survey in August 2019.³

We collected 2,610 responses. Of these, about 20.9% of the responses were from non-managers and about 33.8% were from businesses with less than 50 employees. We exclude those as well as those who indicated that they did not devote full attention to answering the questions (about 9.9%). We also dropped responses from those who finished the survey in an unreasonably short time, i.e., the first percentile of response time. Applying these restrictions, we end up with 1,245 managers. The average response time in this sample was about 7.3 minutes.⁴

In Table 1 we present the summary statistics of the main variables in our survey. The first five variables indicate the share in the control group and each of the four treatment groups. When we launched the survey, we designated each treatment to be randomized evenly across each

³ A growing literature in economics has relied on online survey companies, such as SurveyMonkey and Amazon Mechanical Turk, to conduct online surveys and experiments. Though the respondents collected through these companies are not necessarily representative samples of the population, they do offer a sample that is not too different from the general population, and, as in our case, the possibility to target a specific subset of the population.

⁴ In Appendix Tables 2 and 3 we compare some basic characteristics of our sample relative to the samples in recent papers (Kuziemko et al. 2015, Di Tella and Rodrik 2019) that have used Amazon Mechanical Turk, as well as the American Community Survey (ACS). While our sample is a subset of managers of businesses with 50 or more employees, and employed in the three broad industry sectors, the other samples in Appendix Tables 2 and 3 do not have any explicit restrictions. Appendix Table 2 presents the distribution across states in the US and shows that the geographical distribution of managers in our sample is not very different from that of the other papers, or the ACS. Appendix Table 3 presents the gender, education, racial distribution. The managers in our sample tend to include a higher representation of females than in the overall population. Only a third of our respondents are male. However, the female share is considerably higher in Kuziemko et al. 2015 and Di Tella et al. 2019 as well. Given our focus on managers, the educational attainment of our respondents tends to be higher than in the other samples. In terms of race, our sample of managers have a relatively higher share of blacks and a lower share of whites compared to the other samples.

group, and the resulting distribution reflects this well with each group consisting of approximately 20% of the total sample. In terms of industry, about 42.5% are in healthcare, 38.9% in retail and wholesale, and 18.6% in automotive. The next set of variables are the key outcome variables. In terms of adoption, we ask in how many business processes they would adopt any of the AI technologies in the following year. Respondents were allowed to choose from 0 to 10 or more (i.e., top-coded at 10). On average managers in our sample said that they would adopt AI in about 3.4 business processes.

In terms of AI budget, we ask how much they would budget for AI adoption in dollars, and how they would distribute that budget across the six categories.⁵ The average log AI budget in dollars was 9.45. On average managers allocated 22.4% of the AI budget to R&D, 18.8% to hiring, 16.3% to training, 15% to purchasing AI packages, 12.9% to computing and data resources, and 14.6% to developing AI strategy.⁶ In addition to the R&D budget allocation, we directly ask how they would adjust their workplaces' AI-related innovation activities on a 5-point Likert scale ranging from 1 to 5 (decrease greatly=1, decrease slightly, the same, increase slightly, increase greatly=5).

We examine the five ethical issues when adopting AI, also on a 5-point Likert scale (not important=1, slightly important, moderately important, important, very important=5). On average managers considered each ethical issue more than moderately important, and considered privacy and data security issues the most important. Lastly, we examine how managers would adjust the number of the different types of workers (managers, technical workers, office workers, sales workers, service workers, and production workers) because of AI adoption in a 5-point

⁵ We randomize how the six categories are presented to each respondent, so that the order of the categories do not affect how the percentages are allocated.

⁶ Some of the respondents allocated 100% of the budget to one category. We tried dropping these individuals in the empirical analysis, but the results remain the same.

Likert scale (decrease greatly=1, decrease slightly, the same, increase slightly, increase greatly=5). On average managers responded that they would slightly increase all types of workers, but the technical workers somewhat more.

The empirical analysis that follows examines how the different types of AI regulation affect manager's decision on AI adoption, AI-related budget and allocation, AI-related innovation activities, importance of ethical issues related to AI adoption, and labor adjustment due to AI adoption.

Before examining the regression results, we examine whether the individual and firm characteristics are balanced across the control and treatment groups. Table 2 presents the mean and standard errors of the variables across the control group and four treatment groups. All variables are dummy variables related to the described character. We examine whether each treatment is significantly different from the control. Table 2 shows that the data is balanced across the different treatment groups and randomization was well done, although there is a higher share of black respondents and a lower share of white respondents for the general AI treatment group. In the empirical analysis, we control for all the characteristics in Table 2 by including individual and firm characteristics as fixed effects in the regressions.

5. Results

5.1 Impact of AI regulation on the adoption of AI technologies

Table 3 examines how AI-related regulation affects managers' intention to adopt AI technologies. Specifically, we ask in how many business processes they would adopt AI technologies. The counts range from 0 to 10 or more. Since respondents' choices are top-coded we present both OLS regression results (Panel A) and Censored Poisson regression results (Panel

B). Column 1 presents the regression results that include the four regulation treatments only. In column 2 we control for firm characteristics by including firm size, firm revenue, industry, and state fixed effects. In column 3 we add individual level controls, i.e., gender, race, education, and age fixed effects. Column 4 additionally controls for an index of the firm's human resource management practice and the respondent's role in the organization. Column 5 adds fixed effects for the largest annual budget previously managed by the respondents. Finally, column 6 includes three dummy variables that indicate whether the business currently uses either natural language processing, computer vision, or machine learning at their workplace. Standard errors clustered at the state-industry level are reported in Table 3, and all the following tables. Overall, the coefficient estimates on the four regulation treatments are quite stable across the different columns and are not significantly different across the columns. This indicates that randomization was successfully done. Figure 2 visually illustrates the main results of AI adoption by plotting the coefficient estimates of each regulation treatment from the fully specified model in Table 3 column (6).

The treatment that describes the Algorithmic Accountability Act, i.e., a more general AI regulation (T1), significantly reduces managers' intent to adopt AI technologies in their business processes. Focusing on the OLS results, the general AI regulation treatment reduces the number of business processes that adopt AI by 0.55 (column 6), which is about 16% of the mean value (3.405). The Censored Poisson regression result in column 6 indicates that the general AI regulation treatment reduces AI adoption by 15.7%.

However, the industry-specific AI regulation treatments (T2abc), which offer different treatments across three broad industries by outlining the approaches of the FDA (T2a)(for healthcare, pharmaceutical, and biotech), NHTSA (T2b)(for automotive, transportation, and

distribution), and FTC (T2c)(for retail and wholesale) do not significantly reduce AI adoption. The coefficient estimates are negative but the magnitudes are smaller compared to that of the general AI regulation treatment. Regulation that is more specific to the industry and involves the existing regulatory agency does not have the same negative effect on AI adoption compared to a broad AI regulation that does not concretely reference the regulatory agency in charge of implementing the regulation.

We examine whether the negative effect of the general AI regulation treatment is specific to AI or more of a reaction to regulation in general. Firms using AI are currently subject to existing common laws and statutory requirements, such as tort law and employment, labor, and civil rights law. We remind managers of this through the existing AI-related regulation treatment (T3). This treatment significantly reduces managers' plans to adopt AI technology as well. The negative treatment effect is greater in magnitude than the general AI regulation treatment, although the two are not statistically different. Reminding managers that using AI technology in their businesses will be subject to existing regulation (and potential lawsuits) deters them from adopting AI technology. We interpret these effects as uncertainty with how existing laws govern AI application, and that when reminded of liability, managers assume an adverse position to further adoption. This also suggests a lack of salience of existing laws and regulations.

Finally, we examine the impact of data privacy regulation (T4). The effects of data privacy regulation are not significant without any control (column 1) or with the firm-level controls (column 2) but gradually becomes larger in magnitude and significant in column 6 with the full set of controls. This suggests that there is heterogeneity in the effect of data privacy regulation. Once the control variables are accounted for, data privacy regulation reduces managers' plans to adopt AI technology as well.

5.2 Impact of AI regulation on AI budget and budget allocation.

Next, we examine how regulation information affects how much budget managers would allocate to AI-related activities at the firm, and the allocation of that budget across six different expense categories. Table 4 presents the results from the specification with the full set of control variables.

Columns 1 and 2 examine how many dollars managers would budget for AI adoption at their company in the following year. Column 1 results indicate that there is no significant effect of any of the regulation treatments on the size of the AI budget. We find that there are clusters of responses at multiple of tens and hundreds, and hence are concerned that, despite asking respondents to write in the dollar amount, some may have responded in thousands of dollars. In column 2, we restrict the sample to those who answered with \$10,000 or more. The impact of the agency-specific AI regulation treatment (T2) is now positive and borderline significant at the 5% level. The magnitude is quite large indicating a treatment effect of about 38%. The coefficient estimate on the general AI regulation treatment (T1) is positive at 0.19 as well, though standard errors are larger. AI regulation seems to encourage managers to allocate more to future AI budget.

Columns 3 to 8 examine how managers would allocate that budget across six expense categories in terms of percentage of the total AI budget. By enforcing the allocation to add to 100 percent, we are able to examine the trade-offs managers choose due to AI regulation. We find that AI regulation significantly increases expenditure on developing AI strategy compatible with the company's business strategy (Column 3). The impact is strongest for the general AI regulation treatment (T1), which increases allocation to AI strategy purposes by 3 percentage points, significant at the 5% level. The agency-specific AI regulation (T2) and existing AI-

related regulation (T3) treatments also increase expenditure on developing AI strategy by 2.2 and 2.7 percentage points. The effects of the latter are significant at the 10% level. However, data privacy regulation (T4) has no effect on the budget allocated to developing AI strategy. The general AI regulation treatment also has a positive impact on increasing the budget allocated for hiring the workforce to manage, operate, and maintain AI systems. The increase in budget for developing AI business strategy is primarily offset by a decrease in the budget for training current employees on how to code and use AI technology, as well as purchasing AI packages from external vendors. Figure 3 visually illustrates these results by plotting the coefficient estimates of each regulation treatment. The main takeaway from Table 4 and Figure 3 is that AI regulation forces businesses to “think” and induce managers to expend more on strategizing but at the cost of developing internal human capital.

5.3 Impact of AI regulation on AI-related innovation activities

Table 5 examines whether exposure to AI regulation information affected managers' intent to adjust AI-related innovation activities in the following year. In particular, we ask how they would adjust the following activities: co-operation on AI-related R&D activities with other institutions, such as, universities, research institutes, other businesses; filing AI-related patents; introduction of an AI-related good, service, or production/delivery method that is new or significantly improved. Since respondents were asked to answer these questions on a 5-point Likert scale ranging from 1 to 5 (decrease greatly=1, decrease slightly, the same, increase slightly, increase greatly=5) we present ordered probit regression results that include the full set of control variables. We find that none of the AI-related regulation treatments significantly affect any of the three innovation-related activities in Table 5. Figure 4 visually illustrates these results.

5.4 Impact of AI regulation on ethical issues related to AI technologies

AI regulation information also increases how importantly managers consider various ethical issues when adopting AI (Table 6). Each regulation treatment increases the importance managers put on safety and accident concerns related to AI-technologies, and the existing AI regulation (T3) and data privacy regulation (T4) treatments significantly increase manager perceptions of the importance of privacy and data security. The agency-specific regulation (T2) also increases manager perceptions of the importance of bias and discrimination, and transparency and explainability. Figure 5 visually illustrates these results. Overall, the coefficient estimates are all positive in Table 6, suggesting a general positive effect of AI-related regulation on manager perceptions of the ethical issues related to AI technology.

When asking managers who they think are primarily responsible for AI-related ethical issues at their firm, our results indicate that firm-managers consider themselves to be primarily responsible for ethical issues related to AI (38.6%), followed by: AI package vendors (20.9%), engineers (17.2%), the government i.e. regulatory agencies (16.9%), and the courts (3.9%). The regulation treatments in general do not significantly affect managers' belief on who should primarily be responsible for AI-related ethical issues. However, we find that the agency-specific AI regulation treatment increases managers' beliefs that the court should be primarily responsible for ethical issues (Appendix Table 4).

Our results show that while managers do not devise AI tools, they generally consider themselves responsible for ethical issues related to their implementation. These results suggest that managers face great uncertainty in how existing laws presently govern the use of AI, as well as in relation to quantifying the potential costs of new regulation. This fits well with our findings

that when faced with increased regulation, managers choose to increase strategizing and hire more managers.

5.5 Impact of AI regulation on AI-related labor adjustment

Finally, in Table 7, we examine how AI regulation information might affect employment. Specifically, we ask how managers would adjust the total number of managers, technical workers, office workers, service workers, sales workers, and production workers because of AI adoption. Figure 6 illustrates these results. Exposure to AI-related regulation, in particular, existing AI-related regulation (T3) and data privacy regulation (T4), induces firms to increase the number of managers. The positive impact of AI regulation on the number of managers is consistent with the previous finding that AI regulation induces firms to “think”, by allocating more budget to AI strategy. We find no consistent nor significant impact of regulation on other types of workers.

6. Heterogeneous impact of regulation by industry and by firm size

6.1 Impact of regulation by industry

In this section, we separate the industry- and agency-specific effects of AI regulation. Table 8 presents results on AI adoption, budget allocation, and innovation activity. Table 9 presents results on the ethical issues and adjustment to labor.

AI Adoption

Table 8 column (1) indicates that the negative impact of regulation on AI adoption is especially pronounced in retail and wholesale. All four treatments have a negative impact on the rate of AI adoption, and the magnitudes of the impacts are large and consistent at about a 23% to

28% reduction compared to the control group. For retail, the use of online ads, consumer profiling, digital marketing, and so on, may at present embody greater uncertainty for how revised regulations are likely to impact existing AI practices and use cases. This uncertainty is reflected across all treatments in retail, and significant at the 1% level in relation to data privacy regulation (T4). To a certain extent, our results seem to reflect the current climate that surrounds online platforms, online retail practices and related data handling and consumer profiling, as well as online usage of targeted ad campaigns in which personalized data and related algorithms are used extensively. The negative impact of the general AI regulation (T1) and the existing AI-related regulation (T3) on AI adoption are similar in the healthcare sector, while the negative impact of data privacy regulation (T4) is no longer significant. For automotive, we find no significant impact of regulation on AI adoption across all treatments. While our sample size is smaller for automotive (18.6%), our results suggest that firms operating in the automotive, transportation, and distribution industries, generally factor in a positive outlook on the future of their operations, despite existing laws as well as the mentioning of new and incoming regulations. This positive sentiment is symptomatic of NHTSA's current regulatory approach of removing unintended barriers to AI adoption and innovation.

AI budget and budget allocation

The results in Panel B indicate that AI regulation increases the AI budget as well as the budget share going to developing AI strategy. For automotive these results are consistent and significant at the 1% level under general (T1), as well as agency-specific (T2b) regulation. For general AI regulation, the budgetary increase in developing AI strategy is offset by AI training for existing employees as well as by a budgetary reduction in computing resources and data for AI systems (significant at the 5% level). Data privacy regulation (T4) also increases automotive

budgeting for AI strategy, which again is offset by AI training for existing employees. For retail, the agency-specific regulation (T2c) also increases the budget share allocated to AI strategy. For healthcare, budgeting under agency-specific regulation (T2a) increases the allocation for computer resources and data for AI systems (significant at the 5% level). Under existing AI regulation and laws (T3) healthcare also factors in a budgetary increase for computing resources and data for AI systems, which results in a reduction of purchasing AI packages and systems from existing vendors. Our results show that when faced with the same regulations, the automotive industry is inclined to focus more on increasing its budgets for strategizing, while the healthcare industry devotes more budget to computing resources and data for AI systems. The corresponding budgetary offsets are seen in decreasing AI training for existing workers, as well as in purchasing AI packages, respectively.

Heterogeneity in our results across the healthcare, automotive and retail industries indicates that regulation information is likely to affect industries and their varying compositions in terms of customer relations, business models, data usage, and applied strategic components differently due to industry-specific characteristics. The automotive industries devote the most funds to AI strategy when faced with new rules and regulations. The substantial focus on AI strategy generally reflects the heated competition that currently exists on the market for autonomous-driving-systems, while the prospect of changing regulations and thus market dynamics, forces companies to adjust their strategies even further.

AI-related innovation activities

When we examine the impact of AI regulation on AI-related innovation activities, we find further differential treatment effects across healthcare, automotive and retail (Panel C). AI-related regulation increases managers' plans to file patents in the healthcare sector, while we find

an increase in magnitude as regulation moves from existing AI-related regulation (T3), to agency-specific AI regulation (T2a) (significant at 1% level), and general AI regulation (T1) (significant at 1% level). These findings suggest that as AI regulation increases in scope, so does healthcare manager's intent to file patents. Retail, on the other hand, responds negatively to regulation. When faced with general AI regulation (T1), managers in retail respond by decreasing their intent to file for AI-related patents, and engage in AI-related product or process innovation. For automotive, we find no significant impact of regulation information across all treatments.

Our findings further suggest that industrial idiosyncrasies are present, which makes varying industries respond differently to the same or similar treatments. For healthcare, filing for patents demonstrates a core component of the industry, as we find the intent to file to increase with the scope of the treatments. For retail, we discover a decrease in the intent to file for patents, which signals that other concerns in terms of factoring in future risks are better met by directing attention elsewhere.

Importance of ethical issues related to AI

On ethical issues, we also see some variation across industries (Table 9 Panel D). For automotive, existing AI-related regulation (T3), has a consistent positive impact on ethical issues across safety and accidents, privacy and data security, as well as transparency and explainability. The healthcare industry is more prone to respond positively when faced with general AI regulation (T1) as well as agency-specific regulation (T2a), which increases attention devoted to safety and accidents (significant at the 1% level). For retail, focus on transparency and explainability is positively affected under agency-specific regulation (T2c), significant at the 5% level. We do however find one negative effect, namely that general AI regulation (T1) decreases

privacy and data security concerns in the retail and wholesale industries. The finding might suggest that when uncertainties in existing laws and regulations are exchanged for a broad regulatory framework, managers in retail reduce their concerns over privacy and data security, as the rules for staying compliant become clearer and can more easily be followed.

Adjustment to labor due to AI

In terms of labor (Panel E), the coefficient estimates of all the treatment effects for managers are positive across industries. Whether it be for AI strategizing or concerns over ethical issues, regulation induces firms to increase the number of managers. Another pattern that we see is that the existing AI-related regulation treatment (T3) tends to increase the number of office workers in the automotive sector, which may be a complementary response to increasing the number of managers to deal with potential litigation issues.

6.2 Impact of regulation by firm size

In Tables 10 and 11 we examine how the impact of AI regulation information differs across small versus large firms. We use an annual revenue of \$10 million as the cut off for small and large firms. Table 10 Panel A results show that the negative impact of AI regulation on AI adoption is primarily found for small firms and is statistically strong. Large firms generally are better situated to internalize the costs of regulation, while small firms are faced with hard trade-offs that consistently imply a general reduction in the number of AI processes across all treatments. This potentially suggests that AI regulation is more likely to reduce innovative activity in small firms. Table 10 Panel B indicates that for small firms, general AI regulation (T1) results in an increase in developing AI strategy (significant at the 1% level), which is offset by decreasing AI training for existing employees. For large businesses, on the other hand, this

means hiring more workers related to business' AI systems, which in turn is offset by investments in computing resources and data for AI systems. In relation to data privacy regulation (T4), we find that small firms increase their AI-related R&D, while large firms decrease their AI-related R&D when faced with regulation. This finding suggests that large firms incur a greater costs in terms of restructuring existing practices when faced with data privacy regulation, which implies greater reliance on existing data in AI-related R&D. Smaller and more agile firms may be less reliant on existing data as an input in R&D, which makes them better able to respond to changing practices and data privacy regulations without incurring large costs. While this opens a window of opportunity for smaller firms, an adverse impact is again seen in relation to providing AI training for existing employees.

Table 11 results indicate that AI regulation increases firms' perceptions of the importance of safety and transparency issues in small firms. AI regulation also induces small firms to hire more managers and office workers. Large firms, when reminded of existing AI-related regulation (T3), increase their perception of privacy and data security issues and intend to hire more managers.

7. Conclusion

Our randomized online survey experiment of over a thousand managers tests how information about actual or future AI regulation affects managers' responses about their firms' behavior. We analyze four treatments, each presenting the respondent with different information about: (1) a new general AI regulation involving a new Algorithmic Accountability Act; (2) industry-specific regulations implemented by the FDA, NHTSA and FTC (respectively); (3) existing legal requirements having de facto regulatory effects on AI through common law

doctrines such as tort law, or current statutes governing matters such as employment discrimination; and (4) data privacy regulation, including new statutes such as the California Consumer Privacy Act.

Our findings shed light on how information about the regulatory environment affecting AI affects managers' reported strategic decisions and willingness to adopt AI-related technologies. First, we find that information about industry and agency-specific AI regulation has a less negative impact on firms' reported rate of AI adoption than does general AI regulation. Firms maintain the level of AI adoption under industry-specific regulation but reduce adoption under more general regulation. The industry-specific focus seems to lower the cost managers associate with regulation. Second, we find that regulation induces firms to "think," which we see as an increase in spending on developing AI strategy and hiring more managers. This comes at the cost of hiring other workers and training existing workers. Third, AI regulation reduces AI adoption in small firms and is more likely to reduce their innovative activity. Larger firms seem better positioned to respond to regulatory requirements and bear the costs of regulation. Fourth, industries across healthcare, automotive, and retail respond differently to AI regulation. While AI regulation generally reduces the rate of adoption of AI technologies, the impact varies considerably by targeted industry.

Key implications for AI regulation

Our findings offer several potential implications for the design and analysis of AI-related regulation. First, where possible, regulators should do their best to adapt regulations to the needs and concerns arising in particular industries. Although policymakers sometimes find compelling rationales for adopting broad-based regulatory responses to major problems such as environmental protection and occupational safety, cross-cutting AI regulation such as the

proposed Algorithmic Accountability Act may have enormously complex effects and make it harder to take potentially important sector characteristics into account.

Second, policymakers will do a better job designing and communicating regulatory requirements if they retain a clear focus on regulatory goals. Given the impact of industry sector and firm size on responses, policymakers would do well to take a meticulous approach to AI regulation across different technological and industry-specific use cases. While the importance of certain legal requirements and policy goals — such as reducing impermissible bias in algorithms, and enhancing data privacy and security — may apply across sectors, specific features of particular sectors may nonetheless require distinctive responses. For example, the use of AI-related technologies in autonomous driving systems must be responsive to a diverse set of parameters that are likely to be different from those relevant to AI deployments across drug discovery or online advertising.

Third, given the level of concern among constituencies and target groups for regulation, policymakers should bear in mind the full range of regulatory tools available in the AI context. These include continued reliance on existing legal requirements with relevance to AI such as tort law and employment discrimination that can be gradually elaborated by courts or administrators. Policymakers should also consider the merits of soft-law governance of AI, as well as the costs and benefits of reliance on AI industry standards.

While our findings confirm that conveying information about AI-related regulation generally entails a slower rate of reported AI adoption, we also find that emphasizing existing laws relevant to AI can exacerbate uncertainty for managers in terms of implementing new AI-based solutions. As AI technologies remain at an early stage of adoption, however, the coming magnitude of AI implementation is likely to continue on an upward trending slope, as companies

increasingly will be required to adopt new AI tools and technologies in order to stay competitive. As the potential costs of broad-based general AI regulation are comparable to the costs of existing laws and statutes, this implies that the adoption of clearer rules and regulations could have a net positive effect on the number of firms that are yet to adopt AI technologies. Re-engineering existing AI solutions can be both costly and time-consuming while removing regulatory and legal uncertainties potentially could enable to-be-adopters through the provision of a clearer set of rules and admitted costs of compliance from the outset of adoption. As our study takes the cost side of the equation into consideration, further studies can provide valuable insights about the actual and perceived benefits that potentially come with new forms of AI regulation.

Concluding observations

The extent, content, and responses to AI regulation will no doubt continue to evolve in the years to come, especially in light of the pace of technological innovation and the public's growing exposure to AI. Specific issues such as auditing requirements for algorithms, constraints of sharing of data, rules of governing explainability, and so on, may be addressed in new regulatory requirements or refined interpretations of existing ones. Given the high stakes and the impact not only of substantive requirements but perceptions of those requirements, the public will benefit from a robust, iterative exchange of ideas and information between regulators and business managers, where each can provide meaningful input to each other. Adopters of AI technologies may not always be fully aware of how their AI algorithms function at a detailed technical level. Furthermore, an algorithm that continuously enhances itself based on the progression of data and inputs can make it difficult to determine who is liable as it evolves. Given the pace of innovation and the possibility of lack of understanding in these issues by

managers, clear explanations and information will help managers make well-informed decisions. Pilot studies of AI applications in different sectors, such as autonomous driving, drug discovery processes, and online advertising, may be an essential intermediate step for understanding the implications related to widespread use of AI. Moreover, such pilot studies could involve public-private partnerships and examine how liability could be shared among developers, insurers, the government, and consumers (Kalra and Paddock 2016).

The question of what kinds of regulations are appropriate and most needed by society will remain intricate. No doubt further research that examines the potential impact of AI regulation will help regulators design appropriate AI regulatory frameworks and consider how to implement and adapt existing laws. As policymakers consider the trade-offs, our results underscore the extent to which business managers are sensitive to the risks and costs associated with the regulation of AI. Their responses can have profound effects on workers, businesses, and consumers in the years to come.

References

- Acemoglu, B. D., & Restrepo, P. (2018). The Race between Man and Machine : Implications of Technology for Growth , Factor Shares , and Employment. *American Economic Review*, 108(6), 1488–1542.
- Acemoglu Daron, Pascual Restrepo (2019a). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*. Forthcoming.
- Acemoglu Daron, Pascual Restrepo (2019b). Artificial intelligence, automation and work. Published in *The Economics of Artificial Intelligence: An Agenda* Edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, The University of Chicago Press, Chicago, 2019.
- Aghion, P., Jones, B., & Jones, C. (2017). Artificial Intelligence and Economic Growth. <https://doi.org/10.3386/w23928>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018a). Economic Policy for Artificial Intelligence. *NBER Working Paper Series*, 19, 139–159. <https://doi.org/10.1086/699935>
- Agrawal, A., Gans, J., & Goldfarb, A. (2018b). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.
- Barfield, Woodrow. Pagallo, Ugo (2018) *Research Handbook on the Law of Artificial Intelligence*. Edward Elgar Publishing. Northampton Massachusetts.
- Bessen, J. (2018). AI AND JOBS: THE ROLE OF DEMAND. *NBER Working Paper Series*. <https://doi.org/10.1590/s1809-98232013000400007>
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies: “Engines of growth?” *Journal of Econometrics*.
- Brynjolfsson, E., Rock, D., & Syverson, C. (2017). Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics, 2018, 1–5. <https://doi.org/10.3386/w24001>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2018). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3346739>
- Brynjolfsson, Erik, Tom Mitchell and Daniel Rock (2018). What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers & Proceedings*.
- Clark, J., & Hadfield, G. K. (2019). *REGULATORY MARKETS FOR AI SAFETY*. Conference

- Paper at ICLR 2019.
- Congress. (2019). Algorithmic Accountability Act 2019, 1–15.
- Cuéllar, M. (2019). A Common Law for the Age of Artificial Intelligence: Incremental Adjudication, Institutions, and Relational Non-Arbitrariness. Working Paper.
- Cuéllar, M., & Huq, A. Z. (2019). Privacy ’ s Political Economy and the State of Machine Learning : An Essay in Honor of Stephen J . Schulhofer (forthcoming 2019 in NYU Annual Survey of American Law) Introduction In September 2002 , Stephen Schulhofer published a short book entitled “ T, 873(2005), 1–23.
- Di Tella, Rafael and Rodrik, Dani, Labor Market Shocks and the Demand for Trade Protection: Evidence from Online Surveys (March 2019). NBER Working Paper No. w25705. Available at SSRN: <https://ssrn.com/abstract=3363443>
- Dixen, Jay, Bryan Hong, and Lynn Wu (2019). The Employment Consequences of Robots: Firm-level Evidence. Working Paper.
- Eichler, H. G., L. G. Baird, R. Barker, B. Bloechl-Daum, F. Borlum-Kristensen, J. Brown, R. Chua, S. Del Signore, U. Dugan, J. Ferguson, S. Garner, W. Goettsch, J. Haigh, P. Honig, A. Hoos, P. Huckle, T. Kondo, Y. Le Cam, H. Leufkens, R. Lim, C. Longson, M. Lumpkin, J. Maraganore, B. O’Rourke, K. Oye, E. Pezalla, F. Pignatti, J. Raine, G. Rasi, T. Salmonson, D. Samaha, S. Schneeweiss, P. D. Siviero, M. Skinner, J. R. Teagarden, T. Tominaga, M. R. Trusheim, S. Tunis, T. F. Unger, S. Vamvakas, and G. Hirsch, “From Adaptive Licensing to Adaptive Pathways: Delivering a Flexible Life-Span Approach to Bring New Drugs to Patients,” *Clinical Pharmacology & Therapeutics*, Vol. 97, No. 3, March 2015, pp. 234–246.
- Fenwick M., Vermeulen E.P.M., Corrales M. (2018) Business and Regulatory Responses to Artificial Intelligence: Dynamic Regulation, Innovation Ecosystems and the Strategic Management of Disruptive Technology. In: Corrales M., Fenwick M., Forgó N. (eds) *Robotics, AI and the Future of Law. Perspectives in Law, Business and Innovation*. Springer, Singapore
- Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., ... Rahwan, I. (2019). Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences*, 116(14), 6531–6539. <https://doi.org/10.1073/pnas.1900949116>

- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2018. "Predictably Unequal? The Effects of Machine Learning on Credit Markets." mimeo.
- Goldfarb, A., & Tucker, C. (2012). Chapter Title : Privacy and Innovation. In *Innovation Policy and the Economy* (Vol. 12, pp. 65–89).
- Goldfarb, A., & Tucker, C. E. (2011). Privacy Regulation and Online Advertising. *Management Science*, 57(1), 57–71. <https://doi.org/10.1287/mnsc.1100.1246>
- Goolsbee, A. (2018). Public Policy In An Ai Economy. NBER Working Paper Series.
- Hoffman, M., Kahn, L. B., & Li, D. (2018). Discretion in Hiring. *Quarterly Journal of Economics*, 765–800. <https://doi.org/10.1093/qje/qjx042>. Advance
- Kalra, Nidhi and Susan M. Paddock (2016) Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?. Santa Monica, CA: RAND Corporation, 2016. https://www.rand.org/pubs/research_reports/RR1478.html.
- Kim, J., & Wagman, L. (2015). Screening incentives and privacy protection in financial markets : a theoretical and empirical analysis. *Rand Journal of Economics*, 46(1), 1–22.
- Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva. 2015. "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments." *American Economic Review*, 105 (4): 1478-1508. DOI: 10.1257/aer.20130360
- Lee, Yong Suk and Jong-Hyun Chung (2019). The Evolving Impact of Robots on Jobs: From Displacement to Augmentation. Stanford mimeo.
- McKinsey (2018). Notes from the AI Frontier: AI Adoption Advances, But Foundational Barriers Remain. <https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain>.
- Miller, A. R., & Tucker, C. E. (2011). Can Health Care Information Technology Save Babies ? *Journal of Political Economy*, 119(2), 289–324.
- U.S Department of Transportation. (2018). Automated Vehicles 3.0 - Preparing for the future of transportation
- Varian, H. (2018). Artificial Intelligence, Economics, And Industrial Organization. NBER Working Paper Series. <https://doi.org/10.1590/s1809-98232013000400007>
- Wallach, W., & Marchant, G. (2018). An Agile Ethical/Legal Model for the International and National Governance of AI and Robotics. *Proceedings of the AIES*, 107(3), 7.

<https://doi.org/10.1109/JPROC.2019.2899422>

Weaver, John Frank (2018) Regulation of Artificial Intelligence in the United States. In Research Handbook on the Law of Artificial Intelligence. Edited by Woodrow Barfield and Ugo Pagello. Edward Elgar Publishing. Northampton Massachusetts.

Webb, Michael (2019). The Impact of Artificial Intelligence on the Labor Market. Stanford mimeo.

Figure 1. Research design

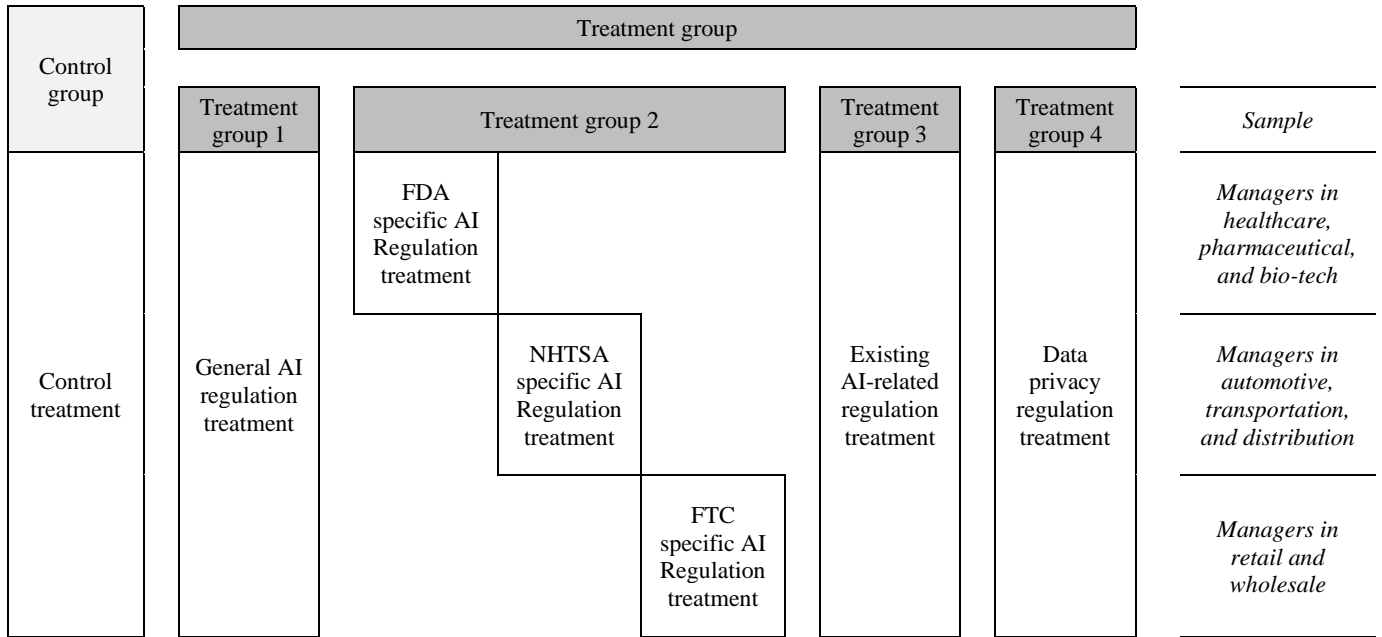
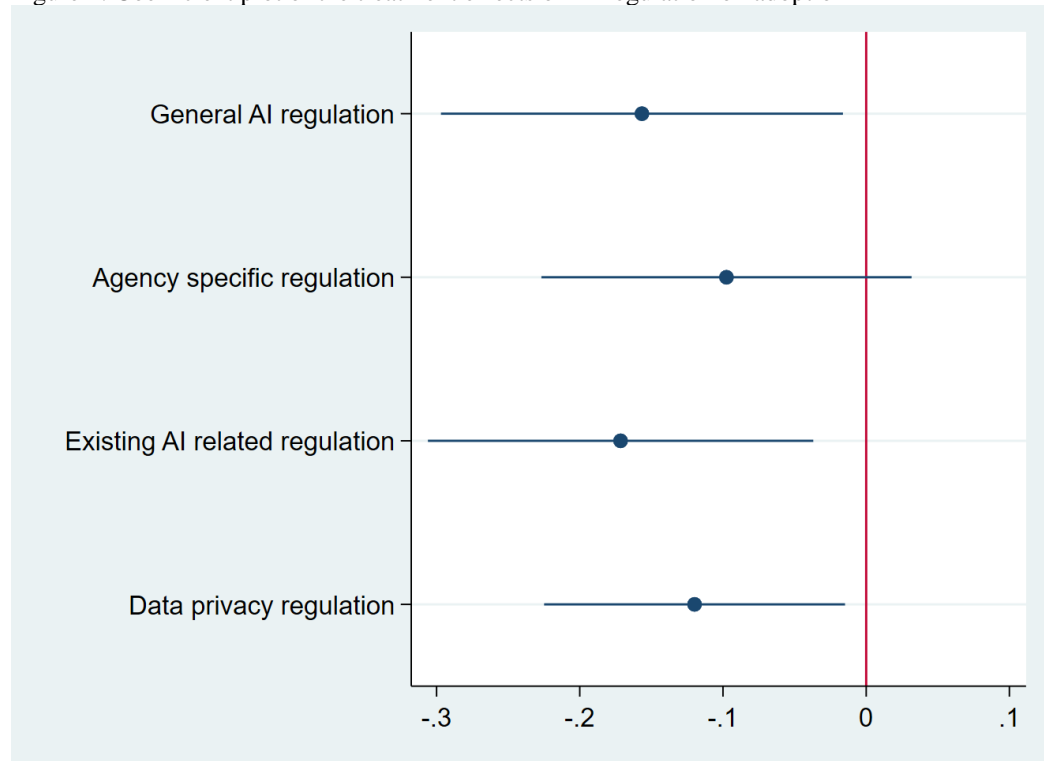
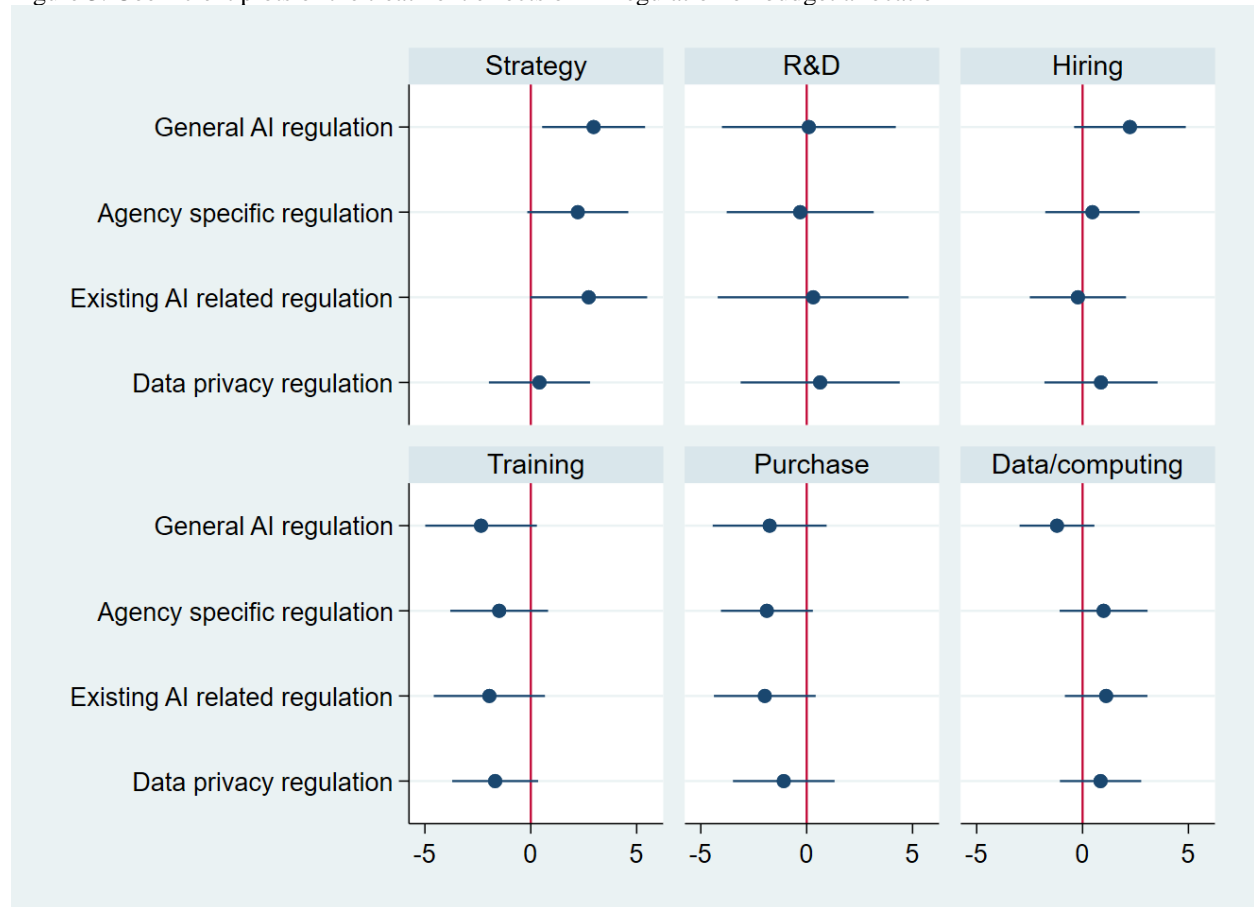


Figure 2. Coefficient plot of the treatment effects of AI regulation on adoption



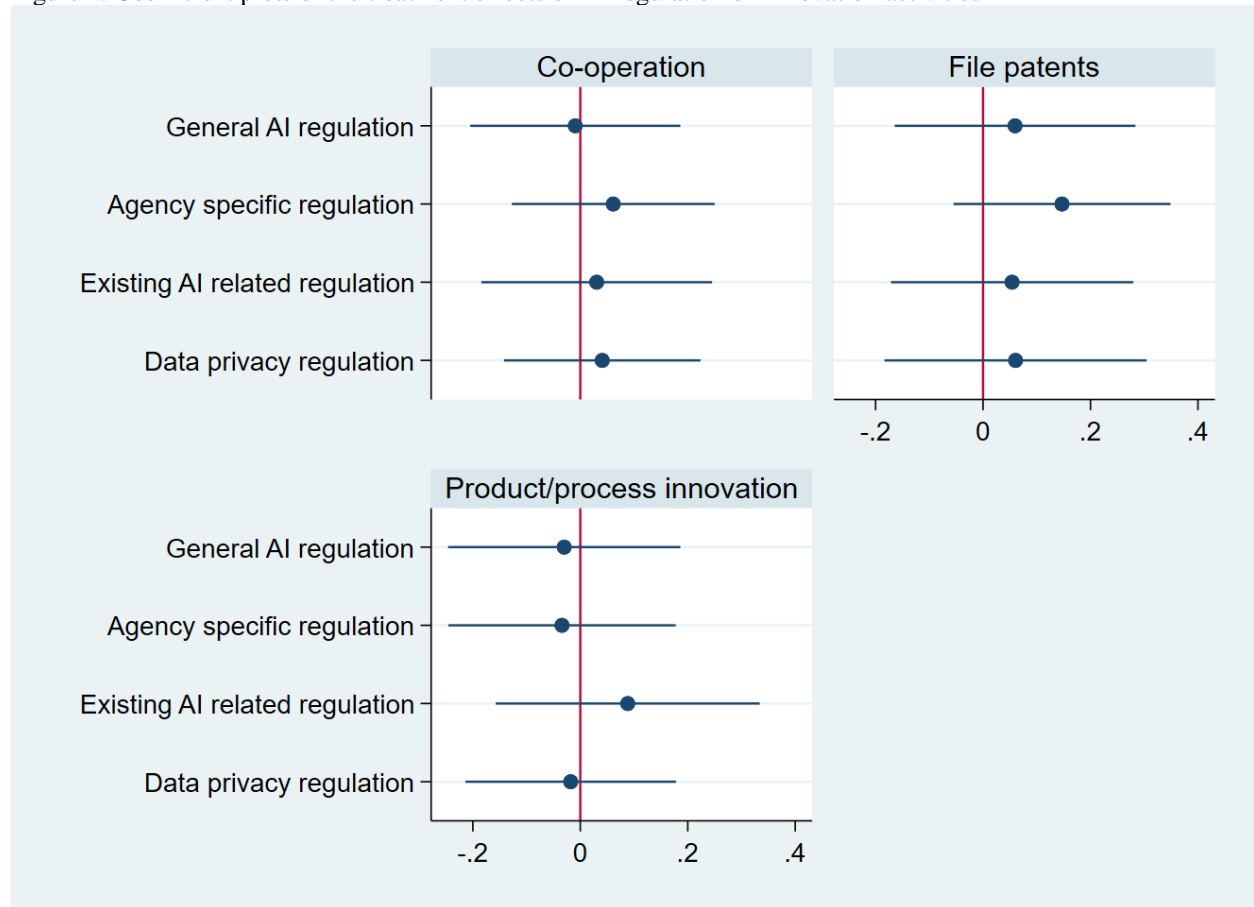
Notes: The dots represent the coefficient estimates from the regression and the bar represents the 95% confidence interval. Each coefficient estimate represents the difference between each treatment group and the control group.

Figure 3. Coefficient plots of the treatment effects of AI regulation on budget allocation



Notes: The dots represent the coefficient estimates from the regression and the bar represents the 95% confidence interval. Each coefficient estimate represents the difference between each treatment group and the control group.

Figure 4. Coefficient plots of the treatment effects of AI regulation on innovation activities




Notes: The dots represent the coefficient estimates from the regression and the bar represents the 95% confidence interval. Each coefficient estimate represents the difference between each treatment group and the control group.

Figure 5. Coefficient plots of the treatment effects of AI regulation on importance of ethical issues



Notes: The dots represent the coefficient estimates from the regression and the bar represents the 95% confidence interval. Each coefficient estimate represents the difference between each treatment group and the control group.

Figure 6. Coefficient plots of the treatment effects of AI regulation on adjustment to labor



Notes: The dots represent the coefficient estimates from the regression and the bar represents the 95% confidence interval. Each coefficient estimate represents the difference between each treatment group and the control group.

Table 1. Summary statistics of key variables

Variable	Mean	Std. Dev.	Min	Max	Obs
Control group	0.194	0.395	0	1	1,245
General AI regulation	0.196	0.397	0	1	1,245
Agency-specific AI regulation	0.214	0.411	0	1	1,245
Existing AI-related regulation	0.204	0.403	0	1	1,245
Data privacy regulation	0.192	0.394	0	1	1,245
Healthcare/pharmaceutical/bio-tech	0.425	0.495	0	1	1,245
Auto/transportation/distribution	0.186	0.390	0	1	1,245
Retail and wholesale	0.389	0.488	0	1	1,245
Number of business processes to adopt AI	3.405	2.777	0	10	1,245
Ln(AI budget)	9.456	4.511	0	23	1,245
Budget share- AI-related research and development	22.393	20.270	0	100	1,245
Budget share-hiring workforce to manage, operate, maintain AI	18.776	14.199	0	100	1,245
Budget share-AI training for existing employees	16.382	12.737	0	100	1,245
Budget share- purchase AI packages from external vendors	14.989	12.260	0	100	1,245
Budget share-computing and data related costs	12.881	11.097	0	100	1,245
Budget share-developing company's AI strategy	14.579	14.948	0	100	1,245
AI innovation activities - co-operation with other institutions	3.714	1.133	1	6	1,245
AI innovation activities - filing patents	3.742	1.170	1	6	1,245
AI innovation activities - produce or process innovation	3.806	1.064	1	6	1,245
Ethical concerns related to AI-layoffs or labor related issues	3.437	1.117	1	5	1,245
Ethical concerns related to AI-racial and gender bias/discrimination	3.461	1.203	1	5	1,245
Ethical concerns related to AI-safety and accidents	3.740	1.103	1	5	1,245
Ethical concerns related to AI-privacy and data security	3.933	1.082	1	5	1,245
Ethical concerns related to AI-transparency and explainability	3.645	1.073	1	5	1,245
Labor adjust from AI adoption-managers	3.370	0.995	1	5	1,201
Labor adjust from AI adoption-technical workers	3.638	0.991	1	5	1,195
Labor adjust from AI adoption-office workers	3.360	1.010	1	5	1,201
Labor adjust from AI adoption-sales workers	3.453	1.037	1	5	1,172
Labor adjust from AI adoption-service workers	3.434	1.041	1	5	1,185
Labor adjust from AI adoption-production workers	3.405	1.013	1	5	1,152

Table 2. Summary statistics of individual and business characteristics by treatment

	Treatment group											
	Control group		General AI regulation		Agency-specific AI regulation		Existing AI-related regulation		Data privacy regulation		Total	
<i>Panel A. Individual characteristics</i>												
Owner or partner	0.166	(0.024)	0.172	(0.024)	0.187	(0.024)	0.118	(0.020)	0.134	(0.022)	0.156	(0.010)
CEO or C-level executive	0.145	(0.023)	0.143	(0.022)	0.135	(0.021)	0.169	(0.024)	0.155	(0.023)	0.149	(0.010)
Managers	0.689	(0.030)	0.684	(0.030)	0.678	(0.029)	0.713	(0.028)	0.711	(0.029)	0.695	(0.013)
Bachelor's degree or above	0.593	(0.032)	0.566	(0.032)	0.547	(0.031)	0.591	(0.031)	0.573	(0.032)	0.573	(0.014)
White	0.664	(0.030)	0.574	(0.032)**	0.622	(0.030)	0.626	(0.030)	0.640	(0.031)	0.625	(0.014)
Black	0.149	(0.023)	0.221	(0.027)**	0.191	(0.024)	0.197	(0.025)	0.163	(0.024)	0.185	(0.011)
Asian	0.054	(0.015)	0.041	(0.013)	0.064	(0.015)	0.043	(0.013)	0.050	(0.014)	0.051	(0.006)
Hispanic	0.075	(0.017)	0.078	(0.017)	0.096	(0.019)	0.098	(0.019)	0.075	(0.016)	0.084	(0.008)
Other	0.021	(0.009)	0.016	(0.008)	0.007	(0.005)	0.008	(0.006)	0.025	(0.010)	0.015	(0.003)
Female	0.656	(0.031)	0.689	(0.030)	0.629	(0.030)	0.650	(0.030)	0.715	(0.029)	0.667	(0.013)
Age less than 30	0.349	(0.031)	0.381	(0.031)	0.348	(0.029)	0.315	(0.029)	0.364	(0.031)	0.351	(0.014)
Age 30 to 45	0.402	(0.032)	0.365	(0.031)	0.419	(0.030)	0.417	(0.031)	0.377	(0.031)	0.397	(0.014)
Age above 45	0.249	(0.028)	0.254	(0.028)	0.232	(0.026)	0.268	(0.028)	0.259	(0.028)	0.252	(0.012)
<i>Panel B. Workplace characteristics</i>												
Small business (less than 500 emp.)	0.456	(0.032)	0.467	(0.032)	0.509	(0.031)	0.433	(0.031)	0.435	(0.032)	0.461	(0.014)
Large business (500 or more emp.)	0.544	(0.032)	0.533	(0.032)	0.491	(0.031)	0.567	(0.031)	0.565	(0.032)	0.539	(0.014)
Revenue less than 1M	0.203	(0.026)	0.262	(0.028)	0.228	(0.026)	0.224	(0.026)	0.201	(0.026)	0.224	(0.012)
Revenue 1M to 9.9M	0.253	(0.028)	0.275	(0.029)	0.281	(0.028)	0.240	(0.027)	0.318	(0.030)	0.273	(0.013)
Revenue 10M to 99M	0.253	(0.028)	0.189	(0.025)*	0.199	(0.024)	0.244	(0.027)	0.234	(0.027)	0.223	(0.012)
Revenue 100M or more	0.290	(0.029)	0.275	(0.029)	0.292	(0.028)	0.291	(0.029)	0.247	(0.028)	0.280	(0.013)
Low management practices	0.481	(0.032)	0.426	(0.032)	0.442	(0.030)	0.437	(0.031)	0.444	(0.032)	0.446	(0.014)
High management practices	0.519	(0.032)	0.574	(0.032)	0.558	(0.030)	0.563	(0.031)	0.556	(0.032)	0.554	(0.014)
Previous budget less than 100K	0.257	(0.028)	0.287	(0.029)	0.262	(0.027)	0.252	(0.027)	0.276	(0.029)	0.267	(0.013)
Previous budget 100K to 999K	0.539	(0.032)	0.500	(0.032)	0.472	(0.031)	0.465	(0.031)	0.464	(0.032)	0.488	(0.014)
Previous budget 1M or more	0.614	(0.031)	0.570	(0.032)	0.607	(0.030)	0.614	(0.031)	0.598	(0.032)	0.601	(0.014)
Natural language processing in use	0.739	(0.028)	0.738	(0.028)	0.734	(0.027)	0.752	(0.027)	0.736	(0.029)	0.740	(0.012)
Computer vision processing in use	0.693	(0.030)	0.717	(0.029)	0.719	(0.028)	0.709	(0.029)	0.745	(0.028)	0.716	(0.013)
Machine learning processing in use	0.763	(0.027)	0.758	(0.027)	0.775	(0.026)	0.752	(0.027)	0.791	(0.026)	0.768	(0.012)
No. of observations	241		244		239		254		267		1245	

Table 3. Adoption of AI

	Number of business processes to adopt AI					
	(1)	(2)	(3)	(4)	(5)	(6)
<u>Panel A. OLS regression results</u>						
General AI regulation	-0.579** (0.235)	-0.524** (0.245)	-0.515** (0.252)	-0.474* (0.254)	-0.513** (0.258)	-0.553** (0.260)
Agency-specific AI regulation	-0.374 (0.244)	-0.298 (0.251)	-0.296 (0.246)	-0.272 (0.243)	-0.325 (0.258)	-0.385 (0.245)
Existing AI-related regulation	-0.511** (0.253)	-0.513** (0.250)	-0.498** (0.250)	-0.489* (0.248)	-0.575** (0.250)	-0.622** (0.246)
Data privacy regulation	-0.295 (0.205)	-0.289 (0.206)	-0.312 (0.196)	-0.308 (0.191)	-0.368* (0.197)	-0.443** (0.196)
Observations	1,245	1,245	1,245	1,245	1,245	1,245
R-squared	0.005	0.061	0.099	0.113	0.157	0.232
<u>Panel B. Censored Poisson regression results</u>						
General AI regulation	-0.167** (0.0679)	-0.150** (0.0689)	-0.147** (0.0704)	-0.136* (0.0706)	-0.152** (0.0709)	-0.157** (0.0716)
Agency-specific AI regulation	-0.105 (0.0682)	-0.0827 (0.0685)	-0.0804 (0.0666)	-0.0770 (0.0654)	-0.0923 (0.0692)	-0.0975 (0.0659)
Existing AI-related regulation	-0.146** (0.0731)	-0.148** (0.0708)	-0.138* (0.0707)	-0.137* (0.0703)	-0.166** (0.0694)	-0.171** (0.0687)
Data privacy regulation	-0.0817 (0.0568)	-0.0816 (0.0563)	-0.0867* (0.0526)	-0.0844* (0.0512)	-0.101* (0.0526)	-0.120** (0.0536)
Observations	1,245	1,245	1,245	1,245	1,245	1,245
Firm level controls	No	Yes	Yes	Yes	Yes	Yes
Individual controls	No	No	Yes	Yes	Yes	Yes
Management controls	No	No	No	Yes	Yes	Yes
Budget experience	No	No	No	No	Yes	Yes
Current AI adoption	No	No	No	No	No	Yes

Notes: Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 4. Budget and allocation

	Budget allocation							
	Log(AI budget)		Developing AI strategy	AI-related R&D	Hiring workers related to business' AI system	AI training for existing employees	Purchase AI package from vendors	Computing resource and data for AI system
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
General AI regulation	-0.0139 (0.421)	0.190 (0.294)	2.966** (1.229)	0.102 (2.076)	2.237* (1.333)	-2.349* (1.333)	-1.749 (1.360)	-1.208 (0.893)
Agency-specific AI regulation	0.506 (0.391)	0.383* (0.197)	2.221* (1.206)	-0.307 (1.754)	0.466 (1.126)	-1.493 (1.168)	-1.880* (1.098)	0.993 (1.049)
Existing AI-related regulation	-0.254 (0.384)	-0.00226 (0.223)	2.735* (1.395)	0.307 (2.279)	-0.221 (1.148)	-1.956 (1.328)	-1.977 (1.214)	1.113 (0.986)
Data privacy regulation	0.198 (0.419)	0.0580 (0.224)	0.410 (1.207)	0.636 (1.899)	0.871 (1.350)	-1.684 (1.025)	-1.083 (1.212)	0.850 (0.971)
Observations	1,245	813	1,245	1,245	1,245	1,245	1,245	1,245
R-squared	0.262	0.347	0.094	0.094	0.084	0.074	0.102	0.080
Firm level controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Management controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Budget experience	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Current AI adoption	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 5. AI-related innovation activities

	Ordered probit regression results		
	Co-operation on AI-related R&D	AI-related patenting	AI-related product or process innovation
	(1)	(2)	(3)
General AI regulation	-0.00666 (0.0919)	0.0550 (0.102)	-0.0307 (0.109)
Agency-specific AI regulation	0.0555 (0.0894)	0.144 (0.0922)	-0.0355 (0.107)
Existing AI-related regulation	0.0276 (0.101)	0.0510 (0.104)	0.0921 (0.125)
Data privacy regulation	0.0407 (0.0866)	0.0563 (0.112)	-0.0178 (0.0988)
Observations	1,245	1,245	1,245
R-squared			
Firm level controls	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes
Management controls	Yes	Yes	Yes
Budget experience	Yes	Yes	Yes
Current AI adoption	Yes	Yes	Yes

Notes: Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 6. Importance of ethical issues related to AI adoption

	Ordered probit regression results				
	Labor issues	Bias and discrimination	Safety and accidents	Privacy and data security	Transparency and explainability
	(1)	(2)	(3)	(4)	(5)
General AI regulation	0.0697 (0.0870)	0.0411 (0.0848)	0.237*** (0.0877)	0.00648 (0.0834)	0.0426 (0.0842)
Agency-specific AI regulation	0.0382 (0.0937)	0.154* (0.0914)	0.300*** (0.0962)	0.0896 (0.103)	0.215** (0.0978)
Existing AI-related regulation	0.0843 (0.111)	0.0112 (0.106)	0.248** (0.102)	0.217** (0.0869)	0.157* (0.0948)
Data privacy regulation	0.146 (0.101)	0.131 (0.105)	0.194** (0.0964)	0.229** (0.109)	0.157 (0.104)
Observations	1,245	1,245	1,245	1,245	1,245
Firm level controls	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes
Management controls	Yes	Yes	Yes	Yes	Yes
Budget experience	Yes	Yes	Yes	Yes	Yes
Current AI adoption	Yes	Yes	Yes	Yes	Yes

Notes: Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 7. Labor adjustment due to AI adoption

	Ordered probit regression results					
	Managers	Technical workers	Office workers	Sales workers	Service workers	Production workers
	(1)	(2)	(3)	(4)	(5)	(6)
General AI regulation	0.134 (0.102)	-0.125 (0.0948)	0.0875 (0.109)	-0.0671 (0.120)	0.0342 (0.112)	0.0180 (0.115)
Agency-specific AI regulation	0.0982 (0.0925)	-0.0474 (0.0907)	-0.0487 (0.0946)	0.0223 (0.0875)	-0.0470 (0.111)	-0.0532 (0.101)
Existing AI-related regulation	0.238** (0.103)	0.0791 (0.0927)	0.0646 (0.100)	0.0577 (0.0896)	0.0270 (0.0956)	0.101 (0.114)
Data privacy regulation	0.209** (0.104)	-0.00362 (0.0923)	0.0153 (0.103)	-0.0569 (0.0862)	0.0315 (0.105)	-0.0455 (0.114)
Observations	1,201	1,195	1,201	1,172	1,185	1,152
Firm level controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Management controls	Yes	Yes	Yes	Yes	Yes	Yes
Budget experience	Yes	Yes	Yes	Yes	Yes	Yes
Current AI adoption	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 8. Impact of AI regulation on adoption, budget allocation, and innovation activity by industry

	A. Adoption	B. Budget allocation							C. Innovation activity		
	No. of business processes to adopt AI	Log(AI budget)	Developing AI strategy	AI-related R&D	Hiring related to business' AI system	AI training for existing employees	Purchase AI package from vendors	Computing resource and data	Co-operation on R&D	AI-related patenting	AI-related product or process innovation
	Censored Poisson (1)	OLS (2)	OLS (3)	OLS (4)	OLS (5)	OLS (6)	OLS (7)	OLS (8)	Ordered Probit (9)	Ordered Probit (10)	Ordered Probit (11)
General AI regulation											
x Healthcare	-0.178* (0.107)	0.0520 (0.664)	1.785 (1.857)	-1.003 (3.295)	3.102 (2.016)	-1.470 (1.943)	-1.570 (2.177)	-0.844 (1.401)	0.0199 (0.127)	0.416*** (0.152)	0.260 (0.163)
x Transportation	0.0631 (0.141)	0.371 (0.774)	7.739*** (2.437)	2.883 (3.737)	1.904 (3.108)	-6.873** (3.091)	-1.716 (3.087)	-3.937** (1.756)	-0.213 (0.211)	-0.233 (0.219)	-0.136 (0.230)
x Retail and wholesale	-0.233* (0.122)	-0.358 (0.686)	2.042 (2.016)	0.564 (3.268)	1.354 (2.189)	-1.339 (1.680)	-2.046 (1.845)	-0.575 (1.449)	0.0679 (0.158)	-0.259** (0.123)	-0.365** (0.148)
Agency-specific AI regulation											
x Healthcare	-0.0336 (0.0947)	1.287** (0.569)	-1.051 (1.758)	-0.453 (2.647)	1.762 (2.003)	-2.157 (1.516)	-2.106 (1.346)	4.005** (1.666)	0.131 (0.146)	0.357*** (0.128)	0.106 (0.154)
x Transportation	0.0508 (0.155)	0.647 (0.730)	6.838*** (2.501)	0.650 (2.512)	-0.571 (2.290)	-2.472 (3.059)	-2.409 (1.834)	-2.036 (1.834)	-0.0310 (0.176)	-0.0593 (0.152)	0.0887 (0.207)
x Retail and wholesale	-0.240** (0.119)	-0.515 (0.597)	3.648** (1.735)	-0.408 (3.321)	-0.433 (1.660)	-0.464 (1.704)	-1.282 (1.761)	-1.061 (1.642)	0.0193 (0.149)	-0.0320 (0.169)	-0.303* (0.172)
Existing AI-related regulation											
x Healthcare	-0.163 (0.102)	-0.231 (0.715)	1.054 (2.045)	0.284 (3.598)	0.559 (1.612)	-1.990 (2.226)	-2.526* (1.522)	2.618* (1.495)	0.0292 (0.130)	0.288* (0.165)	0.230 (0.170)
x Transportation	0.0494 (0.184)	0.399 (0.747)	7.360* (3.838)	-2.642 (4.219)	-1.483 (2.491)	-2.237 (3.292)	-1.061 (3.207)	0.0643 (2.618)	-0.190 (0.205)	-0.0768 (0.155)	0.278 (0.209)
x Retail and wholesale	-0.282** (0.111)	-0.695 (0.568)	2.441 (1.930)	1.975 (3.958)	-0.539 (2.070)	-1.807 (1.690)	-1.872 (2.128)	-0.198 (1.250)	0.130 (0.196)	-0.176 (0.152)	-0.184 (0.211)
Data privacy regulation											
x Healthcare	-0.0941 (0.0826)	0.449 (0.676)	-1.738 (1.924)	0.158 (2.741)	0.245 (2.121)	0.00918 (1.437)	-1.394 (1.481)	2.719* (1.383)	0.0225 (0.137)	0.229 (0.173)	-0.0188 (0.157)
x Transportation	0.139 (0.121)	0.426 (0.725)	6.707* (3.786)	-2.284 (3.207)	1.806 (3.239)	-5.989** (2.839)	0.923 (3.550)	-1.163 (2.068)	-0.166 (0.202)	-0.218 (0.269)	0.119 (0.180)
x Retail and wholesale	-0.263*** (0.0857)	-0.275 (0.709)	-0.123 (1.166)	2.708 (3.643)	0.915 (2.023)	-1.438 (1.596)	-1.657 (2.133)	-0.405 (1.584)	0.155 (0.126)	-0.0530 (0.167)	-0.147 (0.152)
Observations		1,245	1,245	1,245	1,245	1,245	1,245	1,245	1,245	1,245	1,245
R-squared		0.266	0.101	0.097	0.086	0.079	0.103	0.088			

Notes: All regressions include firm level, individual level, management, budget, and current AI use controls. Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Number of observations in the regressions is 1,245. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 9. Impact of AI regulation on the importance ethical issues and adjustment to labor by industry

	D. Importance of ethical issues					E. Adjustment to labor					
	Labor issues	Bias and discrimination	Safety and accidents	Privacy and data security	Transparency and explainability	Managers	Technical workers	Office workers	Sales workers	Service workers	Production workers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
General AI regulation											
x Healthcare	0.0956 (0.105)	0.183 (0.127)	0.326*** (0.115)	0.226 (0.146)	0.0690 (0.137)	0.191 (0.155)	-0.0701 (0.150)	0.0990 (0.160)	-0.228 (0.191)	0.0705 (0.136)	-0.130 (0.164)
x Transportation	0.208 (0.235)	-0.00237 (0.241)	0.343 (0.255)	-0.0222 (0.210)	0.146 (0.182)	0.0799 (0.213)	-0.0718 (0.213)	0.185 (0.172)	-0.123 (0.229)	0.206 (0.305)	0.407 (0.253)
x Retail and wholesale	-0.0199 (0.161)	-0.114 (0.127)	0.0756 (0.165)	-0.246** (0.122)	-0.0250 (0.145)	0.0767 (0.161)	-0.205 (0.165)	0.0159 (0.198)	0.143 (0.191)	-0.0966 (0.210)	0.0286 (0.192)
Agency-specific AI regulation											
x Healthcare	0.0917 (0.145)	0.249* (0.138)	0.307*** (0.111)	0.154 (0.160)	0.175 (0.167)	0.109 (0.139)	-0.0576 (0.136)	-0.0772 (0.139)	-0.0531 (0.126)	-0.0265 (0.155)	-0.152 (0.154)
x Transportation	0.114 (0.186)	0.245 (0.222)	0.568** (0.221)	0.189 (0.238)	0.228 (0.202)	0.185 (0.257)	-0.251* (0.145)	0.117 (0.219)	0.0776 (0.233)	0.0801 (0.218)	0.161 (0.210)
x Retail and wholesale	-0.0533 (0.160)	-0.00882 (0.130)	0.141 (0.187)	-0.0471 (0.155)	0.270** (0.118)	0.0323 (0.126)	0.0939 (0.171)	-0.106 (0.142)	0.0923 (0.127)	-0.140 (0.223)	-0.0288 (0.173)
Existing AI-related regulation											
x Healthcare	-0.0797 (0.175)	0.0156 (0.172)	0.158 (0.156)	0.178 (0.127)	0.0134 (0.156)	0.266* (0.154)	-0.0115 (0.127)	0.0307 (0.132)	-0.101 (0.137)	0.0352 (0.124)	-0.0512 (0.192)
x Transportation	0.124 (0.258)	0.0860 (0.253)	0.450* (0.252)	0.455* (0.240)	0.351* (0.207)	0.360 (0.277)	0.263 (0.207)	0.382** (0.188)	0.235 (0.263)	0.332* (0.184)	0.348* (0.208)
x Retail and wholesale	0.221 (0.164)	-0.0475 (0.152)	0.224 (0.155)	0.121 (0.121)	0.221 (0.135)	0.139 (0.166)	0.0926 (0.175)	-0.0558 (0.202)	0.151 (0.144)	-0.131 (0.190)	0.167 (0.185)
Data privacy regulation											
x Healthcare	0.151 (0.148)	0.0348 (0.173)	0.188 (0.138)	0.155 (0.155)	0.0705 (0.174)	0.112 (0.174)	-0.00270 (0.132)	-0.101 (0.153)	-0.164 (0.141)	-0.0189 (0.179)	-0.200 (0.138)
x Transportation	0.145 (0.254)	0.213 (0.296)	0.195 (0.222)	0.391 (0.307)	0.227 (0.265)	0.442* (0.230)	0.0683 (0.207)	0.276 (0.227)	0.0956 (0.181)	0.196 (0.208)	0.0158 (0.183)
x Retail and wholesale	0.144 (0.166)	0.164 (0.132)	0.178 (0.170)	0.199 (0.179)	0.222 (0.141)	0.185 (0.146)	-0.0228 (0.164)	0.000343 (0.165)	0.00298 (0.148)	-0.00873 (0.161)	0.104 (0.220)
Observations	1,245	1,245	1,245	1,245	1,245	1,201	1,195	1,201	1,172	1,185	1,152

Notes: All regressions include firm level, individual level, management, budget, and current AI use controls. Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Number of observations in the regressions is 1,245. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 10. Impact of AI regulation on adoption, budget allocation, and innovation activity by firm size

	A. Adoption		B. Budget allocation					C. Innovation activity		
	Number of business processes to adopt AI (1)	Developing AI strategy (2)	AI-related R&D (3)	Hiring workers related to business' AI system (4)	AI training for existing employees (5)	Purchase AI package from vendors (6)	Computing resource and data for AI system (7)	Co-operation on AI-related R&D (8)	AI-related patenting (9)	AI-related product or process innovation (10)
General AI regulation										
x Small firm	-0.207** (0.0921)	5.768*** (2.083)	1.144 (2.480)	-1.246 (1.664)	-4.559*** (1.496)	-0.807 (2.067)	-0.301 (1.404)	0.00592 (0.150)	0.159 (0.131)	-0.0280 (0.164)
x Large firm	-0.117 (0.0971)	-0.0249 (1.701)	-0.220 (3.108)	5.379*** (2.050)	-0.0123 (2.066)	-2.818 (1.862)	-2.304* (1.327)	-0.0232 (0.134)	-0.0483 (0.143)	-0.0306 (0.156)
Agency-specific AI regulation										
x Small firm	-0.174** (0.0849)	2.608 (2.200)	4.030 (2.629)	-2.332 (1.634)	-2.136 (1.551)	-2.600 (2.250)	0.431 (1.608)	-0.0337 (0.155)	0.165 (0.143)	-0.0597 (0.148)
x Large firm	-0.0329 (0.103)	2.047 (1.615)	-4.247 (2.656)	2.735 (1.694)	-0.989 (1.679)	-1.135 (1.676)	1.588 (1.412)	0.149 (0.131)	0.132 (0.147)	-0.0121 (0.140)
Existing AI-related regulation										
x Small firm	-0.242** (0.0951)	2.008 (1.884)	1.927 (2.461)	-1.265 (1.675)	-1.698 (1.722)	-1.164 (2.071)	0.192 (1.474)	0.0658 (0.151)	0.0543 (0.140)	-0.0245 (0.168)
x Large firm	-0.109 (0.0935)	3.375* (1.918)	-0.978 (3.011)	0.529 (1.521)	-2.116 (1.798)	-2.703 (1.758)	1.892 (1.426)	-0.00665 (0.140)	0.0447 (0.142)	0.190 (0.144)
Data privacy regulation										
x Small firm	-0.237*** (0.0785)	0.525 (2.018)	6.394** (2.733)	-2.049 (1.698)	-3.024* (1.760)	-2.259 (2.012)	0.413 (1.531)	0.0244 (0.137)	0.170 (0.152)	-0.0262 (0.151)
x Large firm	-0.0237 (0.0846)	0.391 (1.444)	-4.817** (2.257)	3.327 (2.087)	-0.355 (1.682)	0.207 (1.887)	1.247 (1.431)	0.0596 (0.122)	-0.0515 (0.135)	-0.01000 (0.130)
Observations	1,245	1,245	1,245	1,245	1,245	1,245	1,245	1,245	1,245	1,245
R-squared		0.101	0.105	0.115	0.080	0.108	0.084			

Notes: All regressions include firm level, individual level, management, budget, and current AI use controls. Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Number of observations in the regressions is 1,245. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Table 11. Impact of AI regulation on the importance ethical issues and adjustment to labor by firm size (continued)

	D. Importance of ethical issues					E. Adjustment to labor					
	Labor issues	Bias and discrimination	Safety and accidents	Privacy and data security	Transparency and explainability	Managers	Technical workers	Office workers	Sales workers	Service workers	Production workers
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)
General AI regulation											
x Small firm	0.165 (0.132)	0.0810 (0.107)	0.226* (0.126)	-0.0219 (0.115)	-0.0685 (0.104)	0.265* (0.146)	-0.170 (0.144)	0.298** (0.148)	0.164 (0.177)	0.166 (0.161)	0.108 (0.169)
x Large firm	-0.0263 (0.140)	0.00239 (0.153)	0.255* (0.150)	0.0162 (0.146)	0.162 (0.145)	0.0214 (0.136)	-0.0653 (0.137)	-0.105 (0.144)	-0.266* (0.144)	-0.0884 (0.155)	-0.0502 (0.158)
Agency-specific AI regulation											
x Small firm	0.158 (0.144)	0.229* (0.123)	0.385*** (0.123)	-0.00596 (0.150)	0.223* (0.114)	0.252 (0.154)	-0.0144 (0.149)	0.0206 (0.158)	0.268* (0.140)	0.0379 (0.155)	0.0852 (0.137)
x Large firm	-0.0758 (0.132)	0.0762 (0.152)	0.214 (0.134)	0.179 (0.131)	0.201 (0.150)	-0.0409 (0.154)	-0.0787 (0.132)	-0.0796 (0.150)	-0.186 (0.125)	-0.113 (0.167)	-0.176 (0.143)
Existing AI-related regulation											
x Small firm	0.0699 (0.138)	0.0921 (0.132)	0.233 (0.155)	0.136 (0.143)	0.239* (0.140)	0.198 (0.154)	-0.0209 (0.130)	0.180 (0.154)	0.233 (0.151)	-0.0281 (0.167)	0.195 (0.159)
x Large firm	0.0982 (0.149)	-0.0575 (0.147)	0.260* (0.151)	0.285** (0.127)	0.0862 (0.145)	0.283** (0.142)	0.171 (0.162)	-0.0195 (0.143)	-0.0803 (0.143)	0.0836 (0.133)	0.0358 (0.160)
Data privacy regulation											
x Small firm	0.190 (0.141)	0.0671 (0.141)	0.273** (0.127)	0.200 (0.147)	0.228 (0.141)	0.320* (0.170)	0.0115 (0.127)	0.204 (0.143)	0.0663 (0.147)	0.104 (0.152)	-0.0482 (0.165)
x Large firm	0.112 (0.144)	0.211 (0.155)	0.118 (0.142)	0.250 (0.154)	0.0751 (0.151)	0.122 (0.142)	-0.0161 (0.152)	-0.142 (0.145)	-0.131 (0.130)	-0.0188 (0.156)	-0.0103 (0.150)
Observations	1,245	1,245	1,245	1,245	1,245	1,201	1,195	1,201	1,172	1,185	1,152
R-squared											

Notes: All regressions include firm level, individual level, management, budget, and current AI use controls. Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Number of observations in the regressions is 1,245. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.

Appendix Table 1. Treatment texts

Control group	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>While the potential for AI is vast, most organizations still have a long way to go in developing the core practices that enable them to realize the potential value of AI at scale. Business executives and managers will need to think about how to incorporate AI into their business strategy, as well as the transparency and “explainability” of AI algorithms, biases in data, and concerns about safety and privacy.</p>
Treatment 1 – General AI Regulation	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>Until now, states and the federal government have enacted little oversight and regulation specific to AI. But a new Algorithmic Accountability Act is expected to change that. Under this Act, firms that are using or selling AI-related products are subject to a variety of requirements governing their use of AI systems. Requirements include disclosure of firm usage of AI systems, including their development process or contractor of origin, AI system design, model training, and data gathered and in use. The Act also requires firms to disclose to a government agency the impact of their AI systems on safety, accuracy, fairness, bias, discrimination, and privacy. The regulation is expected to go into effect in 2020.</p>
Treatment 2A – Agency-specific AI Regulation (FDA for Healthcare)	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>The healthcare and drug sectors have been actively developing AI technologies for various purposes including patient diagnosis, treatment, drug development, and patient monitoring and care. The Food and Drug Administration (FDA) currently regulates the industry and has proposed a new regulatory framework for AI/Machine Learning-based software. This framework aims to examine and pre-approve the underlying performance of the firm’s AI products before they are marketed, and post-approve any algorithmic modifications. In this process, the FDA will assess the firm’s ability to manage risks associated with various issues such as, transparency and explainability (e.g., diagnosis recommendation algorithms), and security (e.g., use and protection of patient private information) of the AI/Machine Learning based software. FDA’s proposed framework is expected to go into effect in 2020.</p>
Treatment 2B – Agency-specific AI Regulation (NHTSA for Transportation)	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>Autonomous vehicle capabilities have developed rapidly over the last decade and several large companies are currently using cities as testing grounds for unmanned vehicles. The National Highway Traffic and Safety Administration (NHTSA) regulates the autonomous vehicle and logistics industry. NHTSA has specified that its current safety standards constitute an unintended regulatory barrier to innovation of autonomous driving vehicles. For automated driving technologies, NHTSA has emphasized the importance of removing unnecessary barriers and is issuing voluntary guidance rather than regulations that could stifle innovation. NHTSA’s existing regulations and vehicle safety standards remain in effect until a revised framework for automated driving systems is established.</p>
Treatment 2C – Agency-specific AI Regulation (FTC for Retail and Wholesale)	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>The retail sector has been especially fast at deploying and monetizing a range of AI technologies on online and e-commerce platforms. As a result, the Federal Trade Commission (FTC) has engaged in hearings to safeguard consumers from unfair and deceptive practices. For retailers deploying AI technologies, revamped oversight by the FTC will likely require these firms to assess and disclose the impact of their AI systems on various issues. Potential issues include algorithmic discrimination and bias (e.g. in online adds / micro-targeting of consumer groups), transparency (e.g. product recommendation engines) and security (e.g. use and protection of consumers private information). Based on past hearings, new guidelines are expected to be released in 2020.</p>
Treatment 3 – Existing AI-related Regulation	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>Although some observers believe little oversight and regulation has been attached to the area of AI training</p>

	<p>and product deployment, firms using AI technology in the United States generally are subject to common law and statutory requirements. Existing law (e.g., tort law) may require that a company avoid any negligent use of AI to make decisions or provide information that could result in harm to the public. Current employment, labor, and civil rights laws create the risk that a company using AI to make hiring or termination decisions could face liability for its decisions involving human resources. These legal requirements apply now, and will likely continue applying to future products, services, and company practices.</p>
<p>Treatment 4 – Data Privacy Regulation</p>	<p>Recent research has found that early adopters of AI have started to reap the benefits of their investments in this technology. First-movers have already deployed and marketed AI-related solutions across healthcare, autonomous driving, retail and so on. Forty-seven percent of companies say they have embedded at least one AI capability in their business processes.</p> <p>As the development of AI-related products requires more data, policymakers and the public are increasingly concerned about data privacy. For example, California’s recently-enacted digital privacy initiative, the California Consumer Privacy Act of 2018 (CCPA), will affect all businesses buying, selling or otherwise trading the “personal information” of California residents — including companies using online-generated data from residents across their products. In order to stay compliant with the regulation, firms must disclose how they use and store personal data, and how they conform with data privacy rules. California’s regulation goes into effect in 2020. Other states are expected to enact similar data privacy regulations in the near future.</p>

Appendix Table 2. Comparison of state of respondents

State	Our sample	DR (2019)	DDL (2017)	ACS 2015
	% of the total			
Alabama	1.69	1.18	1.29	1.51
Alaska	0	0.11	0.05	0.22
Arizona	2.01	2.27	2.46	2.10
Arkansas	1.2	0.74	0.85	0.92
California	9.24	12.07	9.91	12.12
Colorado	1.29	1.64	1.69	1.69
Connecticut	2.01	0.88	0.97	1.14
Delaware	0.48	0.25	0.39	0.30
District of Columbia	0.4	0.16	0.28	0.22
Florida	5.94	10.92	7.08	6.52
Georgia	4.9	3.38	3.41	3.11
Hawaii	0.72	0.07	0.30	0.45
Idaho	0.24	0.42	0.62	0.49
Illinois	4.58	3.75	4.35	4.00
Indiana	2.81	1.53	2.09	2.03
Iowa	0.48	0.63	0.95	0.97
Kansas	1.04	0.72	0.92	0.88
Kentucky	1.69	1.71	1.49	1.38
Louisiana	1.53	1.13	1.17	1.43
Maine	0.72	0.23	0.50	0.43
Maryland	2.25	1.74	1.84	1.88
Massachusetts	2.57	2.30	2.01	2.18
Michigan	3.86	3.03	3.47	3.11
Minnesota	1.2	1.55	1.51	1.70
Mississippi	0.96	0.83	0.70	0.91
Missouri	1.45	1.58	2.13	1.89
Montana	0.24	0.23	0.22	0.33
Nebraska	0.72	0.46	0.65	0.58
Nevada	0.88	0.83	0.89	0.90
New Hampshire	0.08	0.26	0.50	0.43
New Jersey	2.17	2.20	2.44	2.81
New Mexico	0.24	0.56	0.67	0.64
New York	7.87	6.97	5.71	6.29
North Carolina	3.45	3.43	3.92	3.13
North Dakota	0.4	0.16	0.13	0.24
Ohio	5.46	3.43	4.30	3.63
Oklahoma	1.45	0.91	0.97	1.19
Oregon	0.88	1.62	2.03	1.28
Pennsylvania	4.9	4.20	4.72	4.08
Rhode Island	0.16	0.32	0.25	0.34
South Carolina	1.29	1.57	1.39	1.54
South Dakota	0.24	0.19	0.28	0.26
Tennessee	2.89	1.57	2.08	2.06
Texas	6.91	7.76	7.01	8.18
Utah	0.48	0.72	0.82	0.84
Vermont	0.08	0.33	0.23	0.21
Virginia	1.69	2.83	2.93	2.63
Washington	1.37	2.46	2.78	2.24
West Virginia	0.24	0.53	0.54	0.59
Wisconsin	0.56	1.46	1.91	1.81
Wyoming	0.08	0.12	0.13	0.18

Appendix Table 3. Comparison of individual characteristics

	Our sample	Di Tella and Rodrik (2019)	Di Tella, et al. (2017)	Kuziemko, et al. (2015)	WVS 6 th Wave	ACS 2015
Male	33.25%	46.4%	43.8%	42.8%	48.4%	48.6%
Postgraduate degree	24.18%	17.7%	13.3%	12.6%	11.5%	10.2%
Only college degree	48.43%	49.8%	47.4%	40.7%	24.8%	25.7%
No college degree	27.39%	32.6%	39.3%	46.7%	63.7%	64.1%
White	62.73%	73.1%	80.5%	77.8%	69.8%	74.8%
Black	18.47%	8.8%	9.2%	7.6%	10.4%	12.2%
Hispanic	8.35%	5%	6.6%	4.4%	13.4%	15.5%
Asian	5.14%	6.3%	6.8%	7.6%	-	6.2%
Other race	5.31%	6.6%	2.6%	2.6%	-	2.8%

Appendix Table 4. Primarily responsible for ethical issues

	Primarily responsible for ethical issues				
	Managers	Engineers	Vendors	Government	The court
	(1)	(2)	(3)	(4)	(5)
General AI regulation	-0.358 (0.302)	0.187 (0.123)	-0.0391 (0.129)	-0.205 (0.131)	0.0796 (0.157)
Agency-specific AI regulation	-0.227 (0.246)	-0.0587 (0.125)	-0.213 (0.152)	0.0315 (0.110)	0.354** (0.148)
Existing AI-related regulation	0.0398 (0.259)	-0.0573 (0.124)	-0.199 (0.150)	-0.00874 (0.115)	0.213 (0.130)
Data privacy regulation	-0.182 (0.254)	0.00522 (0.116)	0.0206 (0.149)	0.0502 (0.121)	0.0410 (0.172)
Observations	1,245	1,245	1,245	1,245	1,245
Firm level controls	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes
Management controls	Yes	Yes	Yes	Yes	Yes
Budget experience	Yes	Yes	Yes	Yes	Yes
Current AI adoption	Yes	Yes	Yes	Yes	Yes

Notes: Firm level controls include state, industry, firm size, and firm revenue fixed effects. Individual controls include gender, race, education, and age fixed effects. Management controls include management practice variables related to promotion and firing, and organizational role fixed effects. Budget experience includes dummy variables that control for the largest budget previously managed. Current AI adoption includes dummy variables indicating whether the business currently uses natural language processing, computer vision, or machine learning. Standard errors clustered at the state-industry level are presented in parentheses. ***, **, and * denote statistical significant at 1%, 5%, and 10% level.