


RESEARCH ARTICLE

Investigating the politics and content of US State artificial intelligence legislation

Srinivas Parinandi¹, Jesse Crosson², Kai Peterson¹ and Sinan Nadarevic¹

¹University of Colorado at Boulder, Boulder, CO, USA and ²Purdue University, West Lafayette, IN, USA

Corresponding author: Srinivas Parinandi; Email: srinivas.parinandi@colorado.edu

Abstract

The rapid emergence of artificial intelligence (AI) technology and its application by businesses has created a potential need for governmental regulation. While the federal government of the United States has largely sidestepped the issue of crafting law dictating limitations and expectations regarding the use of AI technology, US state legislatures have begun to take the lead in this area. Nonetheless, we know very little about how state legislatures have approached the design, pursuit, and adoption of AI policy and whether traditional political fault lines have manifested themselves in the AI issue area. Here, we gather data on the state-level adoption of AI policy, as well as roll call voting on AI bills (classified on the basis of consumer protection versus economic development), by state legislatures and analyze the political economy of AI legislation. We find that rising unemployment and inflation are negatively associated with a state's AI policymaking. With respect to individual legislator support, we find that liberal lawmakers and Democrats are more likely to support bills establishing consumer protection requirements on AI usage. The results suggest that economic concerns loom large with AI and that traditional political fault lines may be establishing themselves in this area.

Keywords: federalism; legislative politics; public policy; AI

The presence and increasing utilization of artificial intelligence (AI) technology has the potential to transform the global economy, systems of international security, and even person to person interaction. However, in doing so, it undoubtedly creates challenges for governance. Concerns about how to regulate the use of AI technology by businesses, as well as how to manage the implications of the growth of AI on employment, will no doubt occupy the attention of policymakers around the world in the years moving forward.

As one of the world's largest economies and the home to a key wellspring of AI innovation in the Silicon Valley, the United States will arguably be a testing ground for emerging ideas about constructing policies and a regulatory regime centered on addressing the role of AI technology in the everyday life of the public. Given the potential for AI to influence international trade and security, much of the attention on AI policymaking in the United States, particularly in the future, will focus on the federal government. Particularly given its clear influence on interstate commerce, it is uncontroversial to assume that the federal government will bring its vast resources to bear on developing and adopting a comprehensive AI regulatory strategy in some form in the years to come. At present, however, no such plan—or even dominant set of ideas—exists. As one major law firm, Alston and Bird, puts it “there is no comprehensive federal legislation on AI in the United States”¹; this is a point reiterated in *New York Times*: “Washington has largely been hands off on A.I. rules.”²

¹(Felz and Peretti 2022).

²(Sorkin et al. 2023).

Despite federal inaction, US state governments have stepped into this void and developed their own AI policy agendas over the better part of the last five years. State-level AI policymaking has come from regionally and ideologically heterogeneous source states (both California and Mississippi, for example, have pursued AI policymaking attempts), may form the foundation upon which future federal-level AI regulatory policy is built, and arguably represents the best example of current attempts to regulate AI applications in the United States. More broadly, as recent research has underscored, there is new evidence to suggest that the US states indeed function today as “laboratories of democracy,” in generating policy agendas for pursuit at the federal level.³ Thus, analyzing the nature of state-level AI policy, in terms of which states have pursued AI policymaking, why those states pursued AI policymaking, and how the content of AI policy initiatives has differed across the states, can go a long way toward helping us understand how American AI policy will develop in the years ahead.

In an era of hyperpartisanship and ubiquitous campaigning and electoral competition, however, the emergence of AI is likely to do more politically than generate a new policy area. Indeed, as AI grows to affect the lives of everyday Americans and elites alike, parties may well search for features of AI policymaking that could reinforce party brands and contribute to party competition. As a result, understanding the eventual character of US AI regulation will require an account of how features of AI policy will map onto existing political cleavages. Given the multifaceted nature of AI, it is unclear *ex ante* which aspects of AI regulation will garner the most intense political conflict.

In this article, we offer an empirical first step in understanding these political and policymaking dynamics, with a detailed analysis of attempts across the fifty US states to pursue AI legislation since 2018. At the state level, we trace the adoption of AI legislation across the US states and seek to understand whether prominent economic, political, temporal, and diffusion-related factors have explained the rise of state-level AI policymaking. In order to better understand nascent political dynamics, we also classify state-level AI legislation (including legislation that has been adopted as well as legislation that has not) based on whether it deals with consumer protection or economic development. In doing so, we analyze individual legislator roll call voting data on each kind of AI bill, and we investigate whether individual-level factors such as ideology and partisan affiliation influence support for each type of AI legislation. In tandem, the state adoption analysis and the legislator support analysis give us purchase on the factors most associated with both the adoption of an AI policy regime and how the content of AI policymaking attempts can differ based on political characteristics. Ultimately, in terms of adoption, we find evidence of an association between unemployment and a reduced likelihood of states establishing AI policy regimes, suggesting that state governments are concerned that their attempts to develop AI policy may be perceived as being inattentive to human employment levels. We also find a negative result with respect to inflation and AI policy adoption, which we also believe is due to concerns about optics among lawmakers. At the same time, we find that state governments are more likely to establish AI policy regimes if state government is controlled by the Democratic Party, suggesting that AI policy regimes may not necessarily be bipartisan in character. At the legislator level, we find some evidence that liberals and Democrats are more likely than conservatives and Republicans to support consumer-protection-oriented AI policy while non-consumer protection-oriented AI policy (which essentially is geared toward economic development) lacks clear ideological and partisan dynamics. Together, these results suggest that although AI policy is not completely polarized by any means, traditional ideological and partisan preferences might be manifesting themselves regarding consumer protection specifically. Still, economic factors remain important for understanding AI policymaking dynamics. We proceed as follows. First, we review literature on AI and regulation and then discuss the politics of AI. We then summarize state legislative attempts to regulate AI. Thereafter, we introduce our state adoption and legislator voting analyses and discuss their findings. We conclude by summarizing our findings and offering some thoughts on next steps for understanding the politics and policymaking associated with artificial intelligence.

³(Garlick 2023).

Reviewing the emergence of AI regulation

AI is expected to contribute approximately 15 trillion dollars to the global economy by 2030.⁴ More than simply a new industry, however, AI is changing the way economies operate, pushing firms toward further AI utilization to remain competitive.⁵ As such, increased AI and business automation is expected to create an economic boom, increase productivity, and create a desirable labor supply.⁶ AI can also reduce costs and increase safety and quality.⁷ However, some evidence has also begun to suggest that AI has negative effects on the economy and labor. Acemoglu and Restrepo, for instance, provide evidence that AI is associated with lower rates of hiring, resulting in economic stagnation.⁸ Firms are more likely to allocate tasks to AI if such tasks are compatible, which alters the skill requirements for new employees. Although some argue that AI complements labor,⁹ such positive effects are possibly insignificant when compared to AI's displacement potential.¹⁰

Given the fundamental influence many expect AI to have on the economy and society more broadly, AI regulation has gained considerable traction on the global stage. The European Union, the OECD, and the International Organization for Standardization each have proactively initiated policies toward AI regulation focused on outcomes such as safety and security, privacy, transparency, innovation and development, and promotion of human values.¹¹ The US national government has also implemented federal guidelines to balance consumer protection with industry innovation and development,¹² though such measures constitute only base principles for AI implementation. While some research indicates that AI will negatively disrupt economies and labor markets,¹³ others suggest that such forecasts are overstated.¹⁴ At the consumer level, research indicates that consumers benefit from AI, because AI makes labor more efficient, resulting in lower production costs and lower prices.¹⁵ However, algorithmic bias and discrimination, misuse of individual data, and privacy violations highlight AI's potentially negative externalities.¹⁶

Given these fundamental disagreements, there is no consensus today on how restrictive AI regulation should be.¹⁷ Nevertheless, there *is* a consensus on a general need for AI regulation that simultaneously promotes consumer protection and AI innovation and development.¹⁸ For its part, the US federal government has promulgated multiple initiatives promoting fairness and transparency in AI development,¹⁹ but, it has simultaneously maintained a relaxed stance on legal intervention and has not adopted concrete law regarding the regulation of AI. As such, state legislatures find themselves the forefront of AI regulation, and have intervened with AI policies designed to pursue consumer protection and innovation.

Nevertheless, relatively little scholarly research has examined how states have weighed the competing incentives behind AI regulation. Put differently, little is known about why and which states have adopted regulations, or what factors structure political debates over AI regulation in these states. Instead, existing literature has focused on the substance of individual regulations as adopted. Below, we summarize this literature. We then place it into the broader context of policymaking and diffusion in state legislatures. As we underscore, the novelty of AI policy potentially means that we need to employ

⁴(OxfordInsights 2022).

⁵(Wright and Schultz 2018).

⁶(Arntz, Gregory, and Zierahn 2017; Wright and Schultz 2018).

⁷(Autor 2015; Wright and Schultz 2018).

⁸(Acemoglu and Restrepo 2019).

⁹(Autor 2015).

¹⁰(Acemoglu and Restrepo 2019).

¹¹(Carter 2020; de Almeida, dos Santos, and Farias 2021).

¹²(Chae 2020; The White House 2023).

¹³(Acemoglu and Restrepo 2019; Frey and Osborne 2017).

¹⁴(Arntz, Gregory, and Zierahn 2017; Autor 2015).

¹⁵(Wright and Schultz 2018).

¹⁶(Chae 2020; de Almeida, dos Santos, and Farias 2021).

¹⁷(de Almeida, dos Santos, and Farias 2021; Reed 2018; Tutt 2017; Fosch-Villaronga and Heldeweg 2018).

¹⁸(Acemoglu and Restrepo 2020; Chae 2020; de Almeida, dos Santos, and Farias 2021).

¹⁹(de Almeida, dos Santos, and Farias 2021; Pack 2022).

empirical analysis to understand possible political fault lines regarding the design and adoption of AI policy.

AI regulation and modern politics

In spite of the acknowledged potential benefits to the economy, AI's inherent risks to markets and consumers create an urgent need for regulation.²⁰ Although AI provides many economic benefits to firms and individual consumers, its inevitable increase in use and development likely requires carefully crafted legislation that strikes a fine balance between promoting AI innovation with consumer protection. While there is little disagreement that AI should be regulated, there is nevertheless little consensus on how restrictive regulation should be.

Given the potential economic and individual risks associated with AI, consumer protection is frequently the driving force behind legislative regulation.²¹ Yet, actual regulatory regimes vary considerably in balancing consumer protections against other interests. De Almeida, dos Santos and Farias analyze twenty-one different models of AI regulation and offer their own framework for establishing regulation guidelines, with regulation models ranging from restrictive²² to permissive²³ to somewhere in-between.²⁴ Still others argue that existing legal mechanisms—consumer protection laws within banking regulation—may be extended to regulate AI.²⁵

While the federal government has taken a largely hands-off approach to AI regulation, state legislatures have begun implementing laws dealing with AI innovation and consumer protection. In 2011, Nevada became the first state to enact legislation authorizing the use of AI in the form of autonomous vehicles (AVs).²⁶ A handful of states followed suit, and Tennessee enacted legislation prohibiting local governments from banning vehicles using AI. After gathering information on the impacts of AVs, states began proposing legislation that promotes safety, data, and privacy protections for consumers related to AVs. The regulation of autonomous vehicles illustrates a mixed approach to AI regulation whereby innovation is promoted through the authorization of AV usage and by protecting consumers from the safety and privacy risks that AVs pose. More recently, state legislatures are responding to other implications of AI development and the potential threats to consumer safety, branching far beyond autonomous vehicles.

Current examples of legislation and the difficulty of AI politics

Beyond the regulation of AVs, state legislatures have begun proposing bills that establish committees tasked with studying the effects of AI on labor displacement, its development, and its effect on economic growth. For example, in 2019, Delaware adopted legislation requiring its state agencies to strategically plan for and minimize the risk of AI's labor displacement potential.²⁷ Both Alabama and Illinois adopted policies requiring committees and programs to promote AI innovation and its effect on economic development.²⁸ And New Jersey passed legislation tasking its Commissioner of Labor and Workforce Development to study AI's effect on economic growth.²⁹ These policies illustrate how state legislatures are reacting to AI's economic impact.

Yet there is documented evidence that AI affects more than the broader economy, representing a direct threat to consumers. Cases of AI algorithmic bias and discrimination in obtaining employment,

²⁰(de Almeida, dos Santos, and Farias 2021; Holder et al. 2016; Zardoya Jiménez and Amesti Mendizábal 2022; Holder et al. 2016; Pack 2022).

²¹(Zardoya Jiménez and Amesti Mendizábal 2022).

²²(de Almeida, dos Santos, and Farias 2021; Tutt 2017).

²³(Gurkaynak, Yilmaz, and Haksever 2016; Reed 2018).

²⁴(Fosch-Villaronga and Heldeweg 2018).

²⁵(Zardoya Jiménez and Amesti Mendizábal 2022).

²⁶(Gurkaynak, Yilmaz, and Haksever 2016).

²⁷(NCSL 2022).

²⁸(New York Times Editorial Board 2012).

²⁹(NCSL 2022).

insurance, and credit, raise concerns about AI's application³⁰ and have prompted states to react through regulation addressing potential civil liberties violations. Thus, states like Illinois now prohibit employers and creditors from using AI in ways that consider racial traits in predictive analytics for purposes of establishing employment eligibility or creditworthiness.³¹ Colorado prohibits insurers from using algorithms that discriminate based on race, sex, gender, and other traits.³² And Idaho prohibits the use of algorithmic bias in determining sentencing and bail for defendants.³³ These legislative examples illustrate that state legislatures are aware of AI's potential threats to consumers and are reacting through regulation.

These examples all show that states are responding to the emergence of AI through regulation. However, the factors motivating states to pursue AI policy in the first place, and whether this legislation will focus more on economic growth or consumer protection, are decidedly less well understood. What influences state legislators to embrace AI policy centered on consumer protection? Are the factors driving support for consumer protection-oriented AI policy the same as those driving support for economic development-oriented AI policy? Answering this question is crucial to understanding how the AI industry may take root across the United States, as the question suggests that we explain the rise of a variegated and heterogeneous AI policy regime across the US states.

Literature on the political economy of the US states provides possible guidance. Work on the link between ideology and support for regulation³⁴ suggests that greater governmental liberalism might translate into greater support for consumer protection-oriented AI policy. Other work focuses on a potential liberal tilt to policymaking in general, predicated on the idea that liberals are more likely to view government action as a legitimate pathway to addressing societal problems³⁵). Such work suggests that greater governmental liberalism may even translate to greater support for economic development-oriented AI policy.

At the same time, however, a potential rebalancing of how we understand ideological positions on policy could complicate explanation. Conservative ideology, long thought to be anti-regulatory, might be taking a turn, insofar as the Republican Party espouses more anti-corporate positions. This is especially noteworthy given Republicans' turn toward economic nationalism and rebrand as a working class, populist party. It is possible, then, that traditional ideological expectations may fail to explain AI policymaking in the US states, and it is also possible that traditional partisan expectations (to the extent that ideology maps onto party in the US states) may not hold regarding AI policymaking. Understanding the role of ideology and party in explaining AI policymaking choices is essential to help us navigate the political economy of AI regulation today and potentially moving forward.

In sum, then, there is a need to establish basic empirical facts behind the adoption of and support for different forms of AI regulation. The rapid emergence of AI demands attention from Republican and Democratic legislators alike; however, dominant dimensions of conflict or position-taking heuristics do not necessarily map well onto the AI landscape. Thus, as a first step in understanding the dynamics of AI policymaking, we analyze original data on AI policy adoption and roll call behavior in the US states. As we show, behavior on consumer-facing legislation is decidedly distinct from other sorts of policymaking, though the relationship between left-right ideology and AI policy adoption and support is not deterministic.

AI, partisanship, and policy change in the US States

Given the tandem rise in polarization³⁶ and nationalization³⁷) in American politics, understanding much of policymaking in modern American legislatures revolves around understanding partisan

³⁰(O'Neil 2016).

³¹(NCSL 2022).

³²(NCSL 2022).

³³(NCSL 2022).

³⁴(Parinandi 2023)

³⁵(Boehmke and Skinner 2012).

³⁶(Poole and Rosenthal 1997).

³⁷(Grumbach 2022).

cleavages on particular issue areas. However, given the novelty of AI as an area for policymaking and regulation, it is unclear whether and to what extent AI regulation falls cleanly along existing partisan lines, as we describe above. Thus, as a first cut at exploring the issue dynamics of AI policy, we begin by examining party unity on all existing roll call votes, related to AI policy in the US states.

To do so, we gathered data on all AI legislation in the US states, as identified by the by the National Conference of State Legislatures (NCSL). NCSL maintains a database on all state legislation related to artificial intelligence, from 2018 to 2022. The data were last updated on August 26, 2022.³⁸ The data provide information on AI legislation by year, state, bill name, and status, and each entry includes a brief description of the bills' purposes. Additionally, the data include a hyperlink under the bill's name that allows users to access the full language of the bill on a third-party database. We then build on these data both by independently verifying the accuracy of legislative status and by identifying the roll call votes associated with each AI bill. To do so, we searched LegiScan³⁹ for the relevant roll call votes associated with each bill. For this analysis, we were primarily concerned with final-passage roll calls. Together, our search yielded data on over 1,700 votes on 32 total roll calls, from 2019 to 2022.

These data enable us to more concretely examine the extent to which AI policy does or does not map onto existing partisan cleavages. That is, as a first cut, we can compare the relative level of partisan conflict on AI legislation relative other topics. To do so, in addition to our AI-specific roll call data, we compiled all roll call data available from OpenStates.org for the years closest to those in our AI dataset, 2018–2022 (where available). Using these data, we then measured a simple⁴⁰ index for 1) AI bills during this time period, and 2) all other bills during this time period.⁴¹ With this measure, we aim to examine the extent to which our AI roll calls engender Republicans and Democrats to vote more or less differently, compared to roll calls on other issues.

At base, as Figure 1 displays, AI policy does in fact engender little party unity—at least early on. However, even over AI's rather short history, the average AI vote has indeed grown more partisan. By the final year in our data, in fact, that average AI roll call exhibits higher party unity than the average roll call overall. Still, these trends are quite high-level—meaning it is less clear whether and to what extent other factors may confound the relationship between expressed party unity and roll call behavior. We therefore present a series of models of party unity in Table 1, where we regress each roll call vote's party unity score on an indicator for whether the roll call vote was AI-related. In these models, we also include an indicator for unified/divided government, and an indicator for unified/divided legislatures. Across all models, we include state-level fixed effects to hold constant time-invariant confounds at the state level.

As these models indicate, the baseline level of party unity on AI roll calls is somewhat ambiguous and depends upon model specification. What the regressions do indicate, however—particularly in Model 3—is a noteworthy increase in party unity on AI votes over time. Indeed, although early AI policies did not engender high party unity scores, the interaction term indicates that AI policies have grown in party unity at a faster rate than other issues over the same time.

To be clear, we do not have a strong theory about this growth, nor do we mean to make causal claims. Indeed, it is unclear whether specific sorts of AI bills elicit partisan responses; and it is unclear whether one party, or both parties, are driving this rise. However, as a first cut, these models underscore that partisan features of AI policies have changed, even over the short history of the issue area (Figure 2).

If not partisan division, then what?

If partisanship alone provides only a partial explanation for legislative behavior on AI, what might predict support for the legislation? We posit that three main factors should be important in understanding AI policymaking in state politics. First, given that AI is arguably more human-capital-intensive—both in its creation and in its regulation—than other issue areas, we explore whether the

³⁸<https://www.ncsl.org/technology-and-communication/legislation-related-to-artificial-intelligence>.

³⁹www.legiscan.com.

⁴⁰(Rice 1927).

⁴¹The Rice Index is given by $\left| \frac{Dem_{year} - Dem_{may} - |Rep_{year} - Rep_{may}|}{2} \right|$.

Table 1. Party unity on AI and non-AI votes

	<i>Dependent variable</i>		
	Differenced Party Unity		
	(1)	(2)	(3)
AI Roll Call	-0.060 (0.048)	-0.058 (0.048)	-159.648** (80.830)
AI Roll Call*Year			0.079** (0.040)
Unified Government	0.017 (0.012)	0.023* (0.012)	0.017 (0.012)
Unified Legislature	0.032** (0.013)	0.030** (0.013)	0.032** (0.013)
Year	0.007** (0.001)		0.007** (0.001)
Year Fixed Effects		✓	
State Fixed Effects	✓	✓	✓
Observations	43,943	43,943	43,943
Adjusted R ²	0.046	0.047	0.046

Note: The unit of observation is the legislator-vote.
 ** $p < 0.05$, * $p < 0.10$.

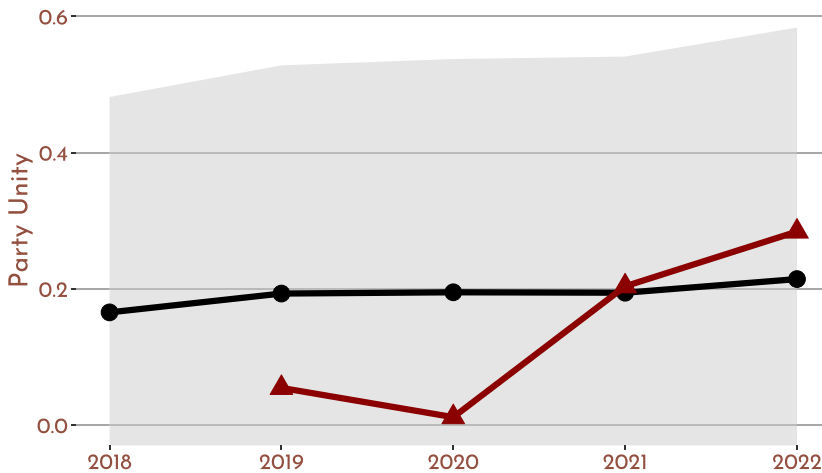


Figure 1. Party unity for AI (red) and all (black) roll calls, 2019–2022.

wealthiness of a state predicts earlier AI policy adoption. Second, as intimated above, previous research highlights that new or novel industries are more likely to experience regulation from left-leaning political bodies, due to general ideological inclinations toward greater regulation of economic entities. Finally, rather than treating AI policy as a single subject or issue area, we posit that policymaking on AI is likely to follow distinct patterns based on the sorts of regulatory factors, namely economic growth and consumer protection, underscored in the aforementioned literature on AI regulation.

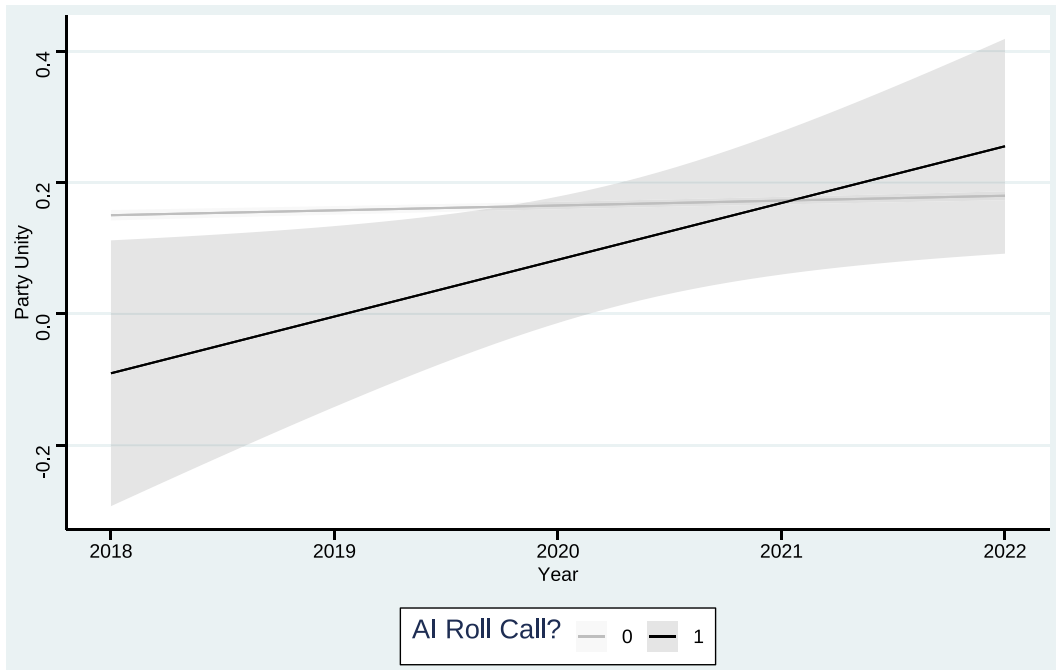


Figure 2. Party unity, AI versus all roll calls (Based on Model 2).

We explore these predictions first by examining the timing and uptake of AI legislation at the state level. Thereafter, we examine individual-level support for AI legislation, making use of our roll call data. Our findings indicate that, although partisan factors like unified government do explain some of the variation in AI policy uptake, economic factors like unemployment seem more predictive. At the individual level, however, party and ideology are indeed quite important for understanding votes on AI issues. Particularly with respect to consumer-focused AI policy, left-leaning and Democratic legislators are more likely to support AI legislation, all else equal.

When and why do states adopt AI policy?

Although roll call votes have shifted in party unity relative to other issue areas, states have undoubtedly also varied on the speed with which—and the basic extent to which—they have pursued AI regulatory policy. In fact, while NCSL identifies 181 AI-relevant bills in US legislatures from 2018 to 2022, only 27 of those bills were enacted. This means that many states have yet to adopt AI-relevant legislation altogether.

To investigate the factors associated with adoption of AI policies (and the timing thereof), we execute a series of adoption models, using event history techniques commonly implemented in the study of policy adoption and diffusion. In doing so, we begin our data frame at the time of the earliest bill introduction, 2017. For every state that has not adopted AI legislation, the dependent variable takes the value 0 in the corresponding year. However, once the state does adopt AI legislation, the variable takes on the value 1 in the year of adoption and the corresponding state exits the dataset thereafter. Structuring the data in this way allows us to examine not only the factors associated with policy adoption, but also the timing of said adoption.

As prefaced earlier, we are interested a series of factors that may explain AI policy adoption at the state level. First, we include *State Per Capita Income*, which enables us to explore whether wealth influences the adoption of AI policy. High household income may offer some degree of protection from threat from AI compared to lower income; wealthier jobs may be less likely to suffer displacement from AI, and even if it occurs, wealthier households may be able to tolerate displacement from AI better than

poorer households. Moreover, this variable captures the possibility that high income individuals who are less likely to possess university degrees (such as plumbers or electricians) might be resistant to displacement by AI and therefore more likely to support the emergence of AI. Along similar lines, we include a battery of other economic variables. *State Unemployment Rate* to some degree captures economic distress in a state and may correspond to public sentiment regarding AI in general: given what we know about loss aversion and how it is potentially heightened given economic difficulty, the public may harbor greater animosity toward the emergence of AI in times of increased unemployment. This animosity (or the belief that it exists) may engender hesitancy to adopt AI policy among legislators given increased unemployment. *State Economic Activity* is a state's share of American economic activity, which captures whether a state's economic size predisposes it to embrace policymaking in the AI space: there is anecdotal support for this, as large states like California and Texas have been epicenters of AI innovation. Finally, *Annual Inflation* captures the yearly US inflation rate and addresses the possibility that rising inflation may lower the mood among policymakers for adopting AI policy, which they fear may be seen as disruptive by voters. Inflation also arguably reduces the tolerance for new policy by businesses, which may be apprehensive about regulatory commitments until uncertainty about prices subsides. This reason would additionally lead to a negative association between inflation and the adoption of AI policy.

In addition to what some may consider classic economic variables, we also include variables designed to measure the human capital development of a state's population—predicated on rival possibilities of how human capital levels may influence AI policy adoption. That is, workforces in states with low human capital might be susceptible to displacement through automation, leading to heightened demand for adopting policy regulating AI in states with low human capital. On the other hand, envisaging how to craft policy dealing with AI may require a citizenry high in human capital, suggesting that states with high levels of human capital development may be more likely to adopt AI policy. The specific variables capturing state human capital levels are: state *Life Expectancy*, and the percentage of a state's population 25 years or older that possesses a *Bachelors Degree or Higher*. We should note that higher education is not a perfect correlate to protection against AI displacement, as some subfields of law have been displaced by AI while other fields of a more service oriented focus (such as plumbing or automotive repair) have not.

We also include a term capturing whether the levers of a state's government are held by the Democrats (*Unified Democratic Government*). Unified government creates conditions more amenable to policy change overall, and inasmuch as AI policymaking constitutes more regulation of industry, it is possible that Democrats will more readily pursue it.⁴² We additionally include a *Governor Party* variable to capture the preferences of a state's chief executive (a value of 1 denotes a Democratic governor). Finally, to address whether left-leaning lawmakers are more eager to craft regulations compared to their right-leaning counterparts, we include the *Average Ideology* of a state's legislature, measured using Shor-McCarty ideal points for state legislators.⁴³

Our final set of variables includes a series of control variables that may explain variation in AI policy uptake. Included among these are *Legislative Professionalism*⁴⁴ and *Neighbor AI Adoption*, which captures whether a state (state i in year t) borders another state that adopted AI policy as of year $t-1$. Both variables are common in models of policy diffusion; and, as they pertain to AI policy may make particularly good sense as predictors. If AI policy does require considerable human capital to understand, for instance, professional legislatures may be better able to pursue AI legislation. Moreover, in the case, of neighbor adoption, policy spillover may be especially beneficial in a technical area like AI. Recognizing that diffusion does not only occur among states that are geographically proximate, we also include a *Ideological Neighbor AI Adoption* variable, capturing whether a state (state j) that is ideologically adjacent to a state (state i) considering to adopt AI policy in year t adopted AI policy as of

⁴²We tried including a variable capturing whether a state possesses a unified legislature, but this variable was dropped due to predicting adoption perfectly.

⁴³(Shor and McCarty 2011).

⁴⁴(Squire 2017).

Table 2. Factors influencing adoption of AI legislation

Variable/Model	(1)	(2)
State Per Capita Income	−0.00 (0.00)	−0.00 (0.00)
Unemployment Rate	−1.04** (0.51)	−0.93** (0.42)
State Economic Activity	0.30* (0.16)	0.21* (0.09)
Annual Inflation	−2.51*** (0.61)	−1.31*** (0.38)
Life Expectancy	0.06 (0.65)	−0.32 (0.38)
Bachelors or Higher	−0.04 (0.13)	−0.01 (0.11)
Average Ideology	−0.13 (1.05)	0.22 (0.96)
Unified Dem Government	2.19* (1.23)	1.79* (0.94)
Governor Party	−0.77 (0.63)	0.01 (0.57)
Legislative Professionalism	3.45 (4.53)	0.47 (3.27)
Neighbor AI Adoption	0.86 (0.94)	−0.20 (0.74)
Ideological Neighbor AI Adoption	−1.95 (1.23)	−0.64 (0.78)
Trump Vote	−0.08 (0.12)	−0.09 (0.08)
Linear Year	2.94*** (0.60)	1.69*** (0.40)
Prior Bills	0.69** (0.32)	0.07 (0.09)
Observations	266	305
Log Likelihood	−36.27	−67.14

***<0.01; **<0.05; and *<0.10 with respect to critical thresholds.

year $t-1$. *Trump Vote* captures the percentage of a state's presidential vote share that was received by Donald Trump in the most recent presidential election before year t and may capture public attitudes against the growth of AI given accounts that the rise of Trump may have been fueled by economic dislocation.

We also include a variable to account for the influence of time (*Linear Year* on adoption, given the over-time dynamics in Table 1. Additionally, we include a variable capturing the number of *Prior Bills* that have been introduced about AI in a given state to account for the possibility that AI policy adoption may occur as a result of previous legislative attempts. Table 2 displays results of logistic estimation with a state's adoption of AI policy as the dependent variable. Model 1 displays results using state-clustered standard errors and an event history structure where a state's opportunity to adopt AI policy disappears once that state has adopted its first AI policy. Model 2 displays results using state-clustered standard errors and an event history structure where a state retains the opportunity to adopt AI policy past its first adoption (that is to say, a state can adopt multiple AI bills into law).

Table 2 displays the results of these regressions. Overall, economic trends exhibit some predictive power for AI policy adoption. First, baseline indicators of economic health are consistent predictors of AI (non)adoption. Both state unemployment and inflation are negatively associated with the adoption of AI policy. As we note above, policymakers may believe that the public views AI technology as a disruptive job destroyer rather than a creator—and therefore hesitate to adopt AI policy in times of economic distress. Beyond economic health, the both models show that states with larger economies (captured by the state economic activity variable) are more likely to adopt AI policy (though at the $p < 0.10$ level). All else equal, governmental attempts to manage the emerging AI sector are more likely to originate in states with the largest economies.

With respect to institutional and partisan variables, unified Democratic government is positively associated with AI policy adoption (though at the $p < 0.1$) level. This provides at least some evidence that Democratic control within state government enables AI policy passage. Conversely, though, governor party does not exhibit a significant correlation with AI policy uptake. The same is true for ideology and a state's vote share for Trump. The findings regarding ideology are noteworthy, given that AI roll calls have grown in divisiveness—and as we show below, are associated with partisanship and ideology at the individual level. At the legislature-level, however, our results suggest that could still be an issue that attracts attention from both sides of the ideological spectrum. Qualitatively, the passage of AI policy in states as ideologically disparate as California and Mississippi indicates that an ideologically heterogeneous group of states have in fact been among the first to develop policy around AI.

Finally, a few contextual factors display significant associations with AI policymaking. The most prominent of these is simply the passage of time: the likelihood of AI policy adoption increases with the linear year variable. Here it is also noteworthy that prior AI bill adoption positively predicts AI policymaking (though this attains significance in only one model). Thus, although party unity on roll call votes has increased in AI policy over time, it has seemingly not hampered states' successful pursuit of policy changes. Other contextual factors, such as neighboring state adoption, were not significant predictors.

Figure 3 displays the influence of rising unemployment on AI policy adoption. Here, continuous control variables are set to their means while binary control variables are set to their most commonly occurring values. Note a negative association that is more evident given increases from low to moderate unemployment (the negative relationship exists but dissipates once unemployment levels of 6 and greater are reached).

Figure 4 displays the influence of rising inflation on AI policy adoption. Again, continuous control variables are set to their means while binary control variables are set to their most commonly occurring values. The negative association between inflation and the probability of AI policy adoption is evident here. Moreover, the association is significant for the majority of the range of the inflation variable (the maximum value of inflation in the dataset is 6.2, suggesting that the effect is significant for nearly the entire range of this variable). Additionally, the steep curve across low to moderate levels of inflation suggests a degree of sensitivity among lawmakers to changes in inflation when considering whether to adopt AI policy.

Readers may be concerned that the adoption-centered approach utilized above misses the true nature of state policymaking on AI since state legislatures can also engage in bill proposal related to AI (the idea here is that focusing only on adoption is akin to looking at the tip of the iceberg regarding state

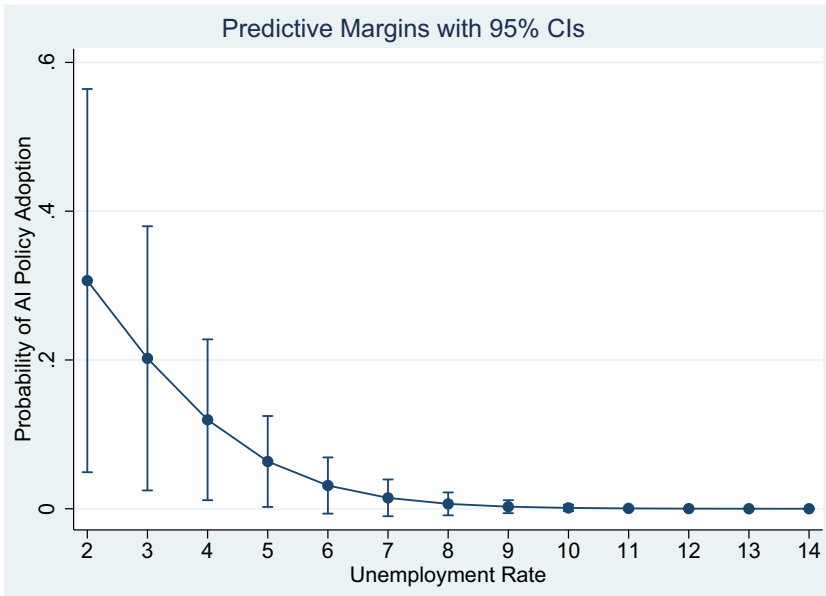


Figure 3. Influence of unemployment on AI policy adoption.

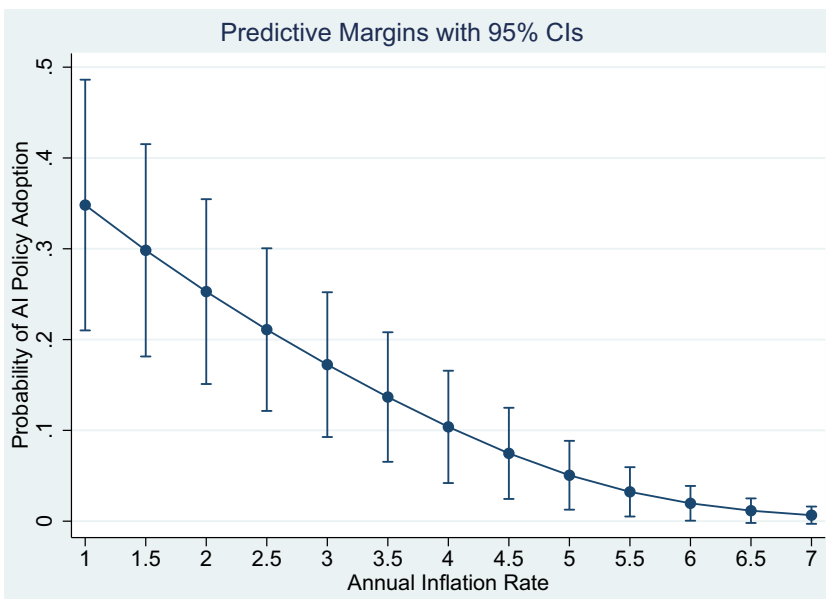


Figure 4. Influence of inflation on AI policy adoption.

policymaking regarding AI). We therefore conduct a robustness check of the adoption analysis displayed above. For this robustness check, we instead use whether a state legislature has proposed a bill related to AI as our dependent variable. Since a state legislature can initiate multiple bill proposals related to AI in a given year, we utilize a conventional time serial cross-sectional data structure where state-year-bill is the unit of analysis (state legislatures that have not proposed an AI bill in a given year receive a value of 0 for that state-year). We utilize the same independent variables and controls as in the earlier adoption analysis except that we now include whether a state legislature has already adopted at least one AI policy (*First Adoption*) or not. Table 3 displays results from this analysis where standard errors are clustered by state.

Table 3. Factors influencing AI bill proposal

Variable/Model	(1)
First Adoption	2.10*** (0.61)
State Per Capita Income	0.00 (0.00)
Unemployment Rate	-0.41*** (0.15)
State Economic Activity	0.01 (0.12)
Annual Inflation	-0.70* (0.36)
Life Expectancy	-0.23 (0.14)
Bachelors or Higher	0.04 (0.11)
Average Ideology	0.42 (0.80)
Unified Dem Government	0.02 (0.79)
Governor Party	0.51 (0.46)
Legislative Professionalism	5.58 (3.41)
Neighbor AI Adoption	0.14 (0.51)
Ideological Neighbor AI Adoption	0.24 (0.54)
Trump Vote	-0.12** (0.05)
Linear Year	1.03** (0.45)
Prior Bills	0.20** (0.09)
Observations	417
Log Likelihood	-114.61

***<0.01; **<0.05; and *<0.10 with respect to critical thresholds.

The analysis focused on bill proposal confirms several of the findings of the adoption-centered approach. As with before, higher unemployment and inflation levels correspond with a reduced likelihood of a state legislature initiating bills related to AI. Similar to before, our argument relies on lawmaker fears about the optics of advancing AI policy: increased unemployment and inflation may signal angst among voters about the economy, and lawmakers may not want to stoke that angst by spearheading AI legislation. Using the bill introduction dependent variable, states that supported Trump's candidacy appear less likely to initiate AI lawmaking attempts compared to states that did not support Trump's candidacy. Moreover, the passage of time and a higher number of previous AI bills introduced in a state increase the probability of a state legislature introducing an AI bill. Unlike the adoption analysis, the Unified Democratic Government variable no longer has statistical support, suggesting that AI bill introduction does not have as strong of a partisan basis as does AI bill adoption. Similar to the adoption results, ideology does not relate statistically with AI bill introduction nor do factors like the partisan affiliation of the governor, legislative professionalism or capacity, whether a geographically adjacent state has adopted at least one policy related to AI, and whether an ideologically adjacent state has adopted at least one policy related to AI. What ultimately stands out from this analysis (and complements the adoption analysis) is that negative economic trends like rising unemployment and rising inflation influence the output of AI policymaking. As an additional way of visualizing the AI policymaking process, we also conduct an analysis where the dependent variable is transformed into an ordered variable capturing how far along an AI bill has progressed in the lawmaking process. Here, we use a dependent variable where 0 captures no AI bill in a given state-year while 1 captures introduction, 2 captures passage in one chamber, 3 captures passage in both chambers, and 4 captures approval of the governor. While we show results of this analysis in the paper's supplemental appendix, the empirical findings largely corroborate the adoption findings reported in Table 2.

While the state adoption analysis showcases the conditions under which states are likely to adopt a first policy dealing with AI (with the results indicating that rising unemployment militates against a state's first adoption of AI policy), the state-level analysis misses important action. We are unable to capture variation within legislature (in terms of ideology, for example) and see how this influences behavior with respect to AI policymaking. The state-level analysis, where we utilized a common event history adoption technique of dropping a state once it adopted its first AI policy, also resulted in our ignoring of multiple AI policy attempts within a state.⁴⁵ Given that legislators may craft multiple AI policy proposals at the same time; and given that any of these may influence subsequent AI policy attempts at the state or federal level, it is important to examine legislator-specific determinants of support for AI policy. Even though ideology does not predict a state's first adoption of AI policy or state's initiation of an AI-related bill, does it predict legislator support for voting "yes" on AI legislation? Are there party-specific effects with respect to analyzing legislator support for AI legislation? In the next section, we divide AI legislation (including bills that have passed as well as those that have not) into two types: those dealing with consumer protection and those dealing with economic growth. This allows us to assess how politics might influence legislator support for advancing different kinds of AI policy.

What predicts support for AI policy changes?

We evaluate the legislator-level features of support for AI legislation by examining the characteristics of "yes" votes on bills dealing with AI. In order to best capture the nature of legislators' support, however, we first classify AI bills according to one major distinction in AI policy, summarized above: consumer-focused versus business-focused legislation. AI technology has brought consumer protection to the forefront and catalyzed questions such as whether (to use a prominent example) resume-screening algorithms will discriminate against applicants with certain names, or whether lending algorithms use a person's listed home address as a reason to *not* extend a mortgage loan. Government action may be necessary to guard against such discriminatory uses of AI technology, and some state legislatures have

⁴⁵(Box-Steffensmeier and Zorn 2002; Jones and Branton 2005).

proposed and even adopted bills to that effect. By contrast, many AI bills do not deal with consumer protection. A large portion of these bills center on economic growth and development. Some state legislatures, for example, have proposed bills claiming that AI technology will transform the economy and introducing means for ensuring that their states will be at the forefront of this economic transformation. We designate a bill as dealing with consumer protection or not after performing a close-reading of the bill. Specifically, we aim to detect language in the bill relating to safeguarding the public from misuse of AI technology and prohibiting the use of such technology to illegally discriminate certain subgroups in the public (for example, by prohibiting life insurance companies from charging higher premiums from life insurance applicants with specific surnames), or we seek to detect language espousing that AI technology should be used to promote equality of opportunity among the public. If either of these conditions are met, we designate a bill as dealing with consumer protection; if neither condition is met, we designate a bill as not dealing with consumer protection. We employed the use of two coders to conduct the classification; they independently classified the bills and then met to resolve inconclusive cases.

Some examples of actual bills provide some context. In Colorado, Senate Bill 113, considered in 2022, establishes rules for the state's government agencies to be transparent with the public about when and how those agencies will use facial recognition technology. Since this bill deals with limits or rules on how AI can be utilized with the public, we consider it to be a bill dealing with consumer protection. In Utah, Senate Bill 96, adopted in 2020, establishes an initiative designed to foster economic development in that state by creating a partnership between Utah's universities and the private sector for the purpose of pursuing collaborative AI endeavors. This piece of legislation does not mention protections, discrimination, or equality with respect to AI and is therefore coded as non-consumer protection oriented. Table 4 details state bill proposals displaying the state name, bill number, type of bill, and bill year.

The table demonstrates that many (though not close to a majority) of states have proposed bills dealing with AI. The states are mainly clustered around the coasts (including California, New Jersey, New York, and Washington); however, there are some states in the interior of the United States (such as Alabama and Idaho) that also developed bills around the issue of AI. Notice also that there is a divide in terms of the topical orientation of the bills. While it may be unsurprising that a conservative state legislature (such as Idaho's or Utah's) may have developed an AI bill that does not center on consumer protection, it is perhaps more surprising that liberal state legislatures (like California's, for example) have authored AI bills that deal with both consumer protection and economic growth. One might think that the party of a legislator explains which of these bills receive a favorable vote to adopt (with conventional wisdom predicting that a Democratic Party affiliation should translate into a "yes" vote on AI bills dealing with consumer protection while a Republican Party affiliation should translate into a "yes" vote on AI bills *not* dealing with consumer protection) but even here, a quick glance of vote returns reveals many legislators not voting based on such a simple expectation. A fuller analysis of the factors motivating legislators to vote "yes" on each kind of AI bill thus is necessary, and we now provide that analysis.

For each of the bills in table 4, we identify roll call votes for legislators and combine these with legislator-specific variables (including ideology, party, chamber, and whether the legislator is a member of a committee with jurisdiction over the bill) and state-specific variables (including the wealth of the state and the unemployment rate of the state). We separate AI bills dealing with consumer protection from those not dealing with consumer protection if the summary or introduction of the bill mentions the rights or need to protect individuals from AI uses in some way; if the summary or introduction of the bill includes language to this effect, we code the bill as dealing with consumer protection (of course, if the summary or introduction do not mention such language, we code the bill as not being centered on consumer protection). We then separately estimate factors predicting a "yes" vote on AI bills dealing with and not dealing with consumer protection. Table 5 displays results for all AI bills, while Table 6 separates bills by type. For each kind of bill, we estimate three models: a logistic model with heteroskedasticity-robust state-clustered standard errors, a multilevel model with state-level random

Table 4. AI bill activity and type

State	Bill Number	Bill Type	Year
California	AB 1576	Consumer Protect	2019
California	ACR 125	Consumer Protect	2019
California	AB 485	Consumer Protect	2019
Illinois	HB 2557	Consumer Protect	2019
California	SB 730	Consumer Protect	2019
Delaware	HCR 7	Consumer Protect	2019
Idaho	HB 119	Consumer Protect	2019
New York	AB 5605	Consumer Protect	2020
Washington	SB 5092	Consumer Protect	2021
Colorado	SB 169	Consumer Protect	2021
Illinois	HB 645	Consumer Protect	2021
Illinois	HB 53	Consumer Protect	2021
Colorado	SB 113	Consumer Protect	2022
Illinois	HB 1811	Consumer Protect	2022
California	SB 1018	Consumer Protect	2022
California	AB 2273	Consumer Protect	2022
Washington	SB 5693	Consumer Protect	2022
California	AB 2408	Consumer Protect	2022
New York	SB 2971	Not Consumer Protect	2019
New York	AB 2946	Not Consumer Protect	2019
Texas	SB 64	Not Consumer Protect	2019
California	AB 594	Not Consumer Protect	2019
California	SJR 6	Not Consumer Protect	2019
California	SB 348	Not Consumer Protect	2019
California	AB 946	Not Consumer Protect	2019
California	SB 444	Not Consumer Protect	2020
New Jersey	SB 2723	Not Consumer Protect	2020
Utah	SB 96	Not Consumer Protect	2020
Alabama	SB 78	Not Consumer Protect	2021
Mississippi	HB 633	Not Consumer Protect	2021
New Jersey	AB 195	Not Consumer Protect	2021
New Jersey	SB 2723	Not Consumer Protect	2021
Illinois	SB 252	Not Consumer Protect	2022

effects, and a multilevel model with both state- and bill-level effects. We do *not* utilize state-specific fixed effects because of the cross-sectional nature of our data, which renders the use of fixed effects inappropriate since state-specific right-hand-side variables are slow-moving and do not change.

Table 6 displays the first set of findings. Pooling all types of bills, and across all model specifications, conservative ideology is negatively associated with voting yes on AI policies. The same is true for

Table 5. Factors associated with yes votes on AI legislation

	<i>Dependent variable</i>		
	Pr("Yes" Vote)		
	(1)	(2)	(3)
Conservative Ideology	−0.994** (0.140)	−0.074** (0.009)	−0.074** (0.009)
Member of Relevant Committee	0.495** (0.232)	0.042** (0.016)	0.037** (0.017)
Upper Chamber	−0.601** (0.152)	−0.039** (0.011)	−0.053** (0.012)
Republican Party	−0.468 (0.307)	−0.050** (0.021)	−0.050** (0.021)
Trump Vote	−0.007 (0.040)	−0.004 (0.010)	−0.003 (0.012)
Democratic Governor	0.502 (0.771)	−0.146 (0.191)	−0.127 (0.244)
Unemployment Rate	0.366** (0.090)	0.026** (0.009)	0.013 (0.011)
Annual Inflation	1.025** (0.238)	0.098** (0.034)	0.066 (0.044)
Life Expectancy	−0.447** (0.120)	−0.048* (0.025)	−0.046 (0.032)
Bachelor's or Higher	0.214** (0.049)	0.005 (0.013)	0.008 (0.018)
Per Capita Income	−0.000** (0.000)	−0.000 (0.000)	−0.000 (0.000)
State Economic Activity	0.161** (0.043)	0.003 (0.014)	0.006 (0.018)
Linear Year	−1.672** (0.433)	−0.192** (0.090)	−0.145 (0.119)
State Random Effects		✓	✓
Bill Random Effects			✓
Observations	3,003	3,003	3,003
Log Likelihood	−702.389	−389.529	−368.974
Akaike Inf. Crit.	1,432.778	811.057	771.948
Bayesian Inf. Crit.		907.175	874.074

Note: The unit of observation is the legislator-vote.

** $p < 0.05$, * $p < 0.10$.

Table 6. Factors associated with yes votes on AI legislation, by type of legislation

	<i>Dependent variable</i>					
	Pr("Yes" Vote)					
	(4)	(5)	(6)	(7)	(8)	(9)
Conservative Ideology	−0.400 (0.264)	−0.017 (0.013)	−0.017 (0.013)	−0.562** (0.215)	−0.085** (0.012)	−0.085** (0.012)
Member of Relevant Committee	0.926 (0.620)	0.032 (0.022)	0.033 (0.022)	0.676** (0.308)	0.044* (0.022)	0.034 (0.022)
Upper Chamber	−0.021 (0.293)	−0.001 (0.016)	−0.004 (0.016)	−1.299** (0.218)	−0.086** (0.015)	−0.097** (0.016)
Republican Party	0.434 (0.520)	0.018 (0.028)	0.018 (0.028)	−2.725** (0.546)	−0.158** (0.029)	−0.157** (0.028)
Trump Vote	0.245** (0.111)	0.015** (0.006)	0.013** (0.007)	2.451* (1.323)	0.069 (0.097)	0.021 (0.068)
Democratic Governor	5.220** (2.132)	0.285** (0.106)	0.255** (0.122)	35.933 (27.425)	0.955 (1.856)	0.136 (1.398)
Unemployment Rate	0.220 (0.236)	0.013 (0.010)	0.010 (0.011)	1.834** (0.702)	0.060 (0.054)	0.028 (0.036)
Annual Inflation	0.007 (0.583)	0.007 (0.027)	0.004 (0.032)	7.451** (1.773)	0.250** (0.095)	0.166** (0.070)
Life Expectancy	0.280 (0.479)	0.018 (0.021)	0.019 (0.025)	−2.020 (1.747)	−0.110 (0.146)	−0.138 (0.102)
Bachelor's or Higher	0.413 (0.289)	0.016 (0.012)	0.011 (0.015)	−0.710* (0.392)	−0.023 (0.022)	0.002 (0.019)
Per Capita Income	−0.000** (0.000)	−0.000** (0.000)	−0.000 (0.000)	0.001** (0.000)	0.000 (0.000)	0.000 (0.000)
State Economic Activity	0.422** (0.155)	0.018** (0.007)	0.016* (0.009)	0.625 (0.427)	0.029 (0.041)	0.020 (0.028)
Linear Year	1.132 (0.885)	0.040 (0.044)	0.042 (0.054)	−19.266** (4.839)	−0.597** (0.230)	−0.399** (0.183)
Consumer Protection Bill?				✓	✓	✓
State Random Effects		✓	✓		✓	✓
Bill Random Effects			✓			✓
Observations	1,293	1,293	1,293	1,710	1,710	1,710
Log Likelihood	−265.630	9.505	9.766	−355.812	−286.668	−267.254
Akaike Inf. Crit.	559.260	12.990	14.468	739.623	605.336	568.509
Bayesian Inf. Crit.		95.626	102.268		692.444	661.061

Note: The unit of observation is the legislator-vote.
 ** $p < 0.05$, * $p < 0.10$.

membership in the Republican Party in the models with state and bill-level effects.⁴⁶ In the multilevel model results specifically, Republicans are approximately 5 percent less likely to vote yes on AI legislation. Moreover, given the relative tight range of the ideology measure (roughly -3 to 3), the -0.074 coefficient on conservatism is also noteworthy substantively. In the model without state- or bill-level effects, economic factors such as per capita income, unemployment rate, and inflation continue to be significant predictors. However, in these cases, the signs are in the *opposite* direction of the adoption models: poor economic conditions seem to be *positively* correlated with voting 'yes' on AI policy. It is worth noting that these variables fluctuate in significance based on specification; however, the difference in results compared to the adoption models is important to underscore.

These differences are better illuminated when examining Table 6, where we run analyses separately on consumer protection and non-consumer protection votes. As the table underscores, nearly all of the partisan and ideological findings in Table 6 are driven by roll calls on consumer protection bills. In the random effects specifications,⁴⁷ the coefficient on party is quite strong for consumer protection bills: Republican legislators are significantly less likely (experiencing a 16 percentage point decrease in the probability of voting "yes") to support consumer protection AI legislation. Of course, proper context is also necessary, as the reduction in probability goes from 99.8 to 84 percent, suggesting (as is well known) that bills going to a final vote on adoption are likely to pass. Nonetheless, the existence of partisan effects suggests that consumer protections in AI are more likely to be supported by Democrats. Even more striking is that this partisan association comes in addition to a noteworthy ideological association: conservative legislators are significantly less likely to vote 'yes' on consumer protection AI measures—while not exhibiting a consistent association with yes votes on non-consumer protection measures.

In addition to findings about partisanship and ideology, our models consistently point to affiliation with a committee with jurisdiction over a consumer protection bill to make a "yes" vote more likely. While the coefficient is smaller than partisanship and ideology, the results do suggest that committee jurisdiction matters, at least in on consumer protection bills. Interestingly, though, the opposite is true for life expectancy: legislators are *less* likely to vote yes on consumer protection bills in higher life expectancy states.

In the models pertaining to non-consumer protection AI bills, results are far more sparse. There are few partisan or ideological associations of note, and even economic indicators are not especially predictive. These findings matter substantively as they suggest that switching to non-consumer protection topical matter does not automatically confer Republican support for AI legislation. It may be that Republicans are somewhat skeptical about AI legislation in general compared to Democrats, though we cannot conclusively say that Republicans are *opposed* as the non-consumer protection AI results are not statistically significant. The same is true for our (non)findings on legislator ideology. One might imagine that there should be an ideological pattern to voting for consumer protection AI bills as well as non-consumer protection AI bills, with liberals voting for consumer protection and conservatives voting for non-consumer protection bills. However, results from our analysis suggest that this is not the case. It is possible that ideological fault lines with respect to AI legislation have yet to fully form, with consumer protection bills engendering more partisan and ideological dynamics more readily than non-consumer protection bills. It is also possible that at least within the Democratic Party coalition, liberals and moderates can agree upon the goal of implementing consumer protection-oriented AI legislation.

One concern is that even though we include bill-specific and state-specific random effects in our analysis, we have not provided estimates that cluster standard errors by bill. We therefore replicate the bill-level analysis using bill-clustered standard errors with three logistic model specifications in Table 7: a model evaluating "yes" votes on consumer protection-oriented AI bills; a model evaluating "yes" votes on non-consumer protection-oriented AI bills; and a model evaluating "yes" votes on all AI bills.

⁴⁶Though, interestingly, Trump vote does not reach statistical significance in any of the models.

⁴⁷Concerning AI bills dealing with consumer protection, we urge paying attention to the state random effects specification, as this is supported by the likelihood ratio finding.

Table 7. Factors associated with “Yes” votes on AI legislation

Variable/Model	Consumer Protection	Non-Consumer Protection	All
Conservative Ideology	−0.674*** (0.252)	−0.399 (0.282)	−0.993*** (0.240)
Relevant Committee	0.789*** (0.290)	0.925* (0.524)	0.495** (0.215)
Upper Chamber	−1.120** (0.465)	−0.021 (0.171)	−0.600** (0.266)
Republican	−2.305*** (0.766)	0.433 (0.545)	−0.467 (0.484)
cv Trump Vote	−0.140 (0.918)	0.245*** (0.051)	−0.006 (0.092)
Unemployment Rate	0.302 (0.517)	0.219* (0.132)	0.365*** (0.109)
Annual Inflation	0.452 (0.316)	0.006 (0.144)	1.025** (0.400)
Life Expectancy	−0.479 (1.390)	0.279 (0.287)	−0.446 (0.294)
Bachelor’s or Higher	0.129 (0.280)	0.412*** (0.119)	0.214* (0.119)
State Per Capita Income	−0.000*** (0.000)	−0.000*** (0.000)	−0.000** (0.000)
State Economic Activity	0.119 (0.327)	0.422*** (0.070)	0.160 (0.119)
Governor Party	0.176 (20.228)	5.220*** (1.130)	0.501 (1.537)
Linear Year	◇	1.131*** (0.335)	−1.671* (0.944)
Observations	1710	1293	3003
Log Likelihood	−373.43	−265.63	−702.38

***<0.01; **<0.05; and *<0.10 with respect to critical thresholds. All standard errors are **clustered by bill**.

◇: Model does not converge with inclusion of linear year variable.

The table provides some degree of corroboration with results from the earlier bill-specific analyses. As before, conservative ideology and Republican partisan affiliation are negatively associated with voting yes on consumer protection-oriented bills, suggesting the emergence of ideological and partisan preferences for consumer protection-oriented AI legislation. As before, there is less evidence of ideological or partisan effects with respect to non-consumer protection-oriented AI bills, reiterating the idea that ideological or partisan fault lines may not have formed with respect to non-consumer protection-oriented legislation. The economic variables that predicted “yes” voting on consumer protection-oriented bills in the previous analysis, namely unemployment rate and annual inflation, are now non-significant with respect to consumer protection bills (though they are still positive and inflation almost attains statistical significance). These variables retain their significance and

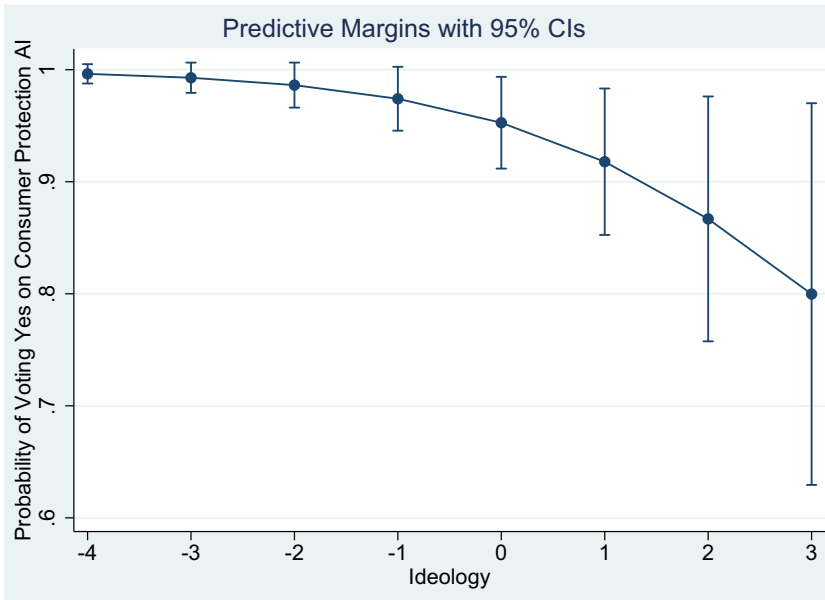


Figure 5. Influence of ideology on consumer protection AI voting.

directionality with respect to all AI legislation, however, suggesting that legislators do support AI measures in times of economic distress even if those same legislators are less likely to adopt AI policy during those same distressing times. More study is necessary to explore this, but what is possible is that legislators see the value of AI legislation as an economic development tool and (if the state and bill-specific random effects results from the previous table are believed) consumer protection tool, but delay adopting AI policy until times of economic distress have dissipated.

Figures 5 and 6 respectively display the influence of increases in conservative ideology and Republican Party affiliation (versus a baseline of Democratic affiliation) on the probability of voting yes on consumer protection-oriented AI legislation. Note that increases in conservatism (especially from a neutral ideological score of 0 to the right) correspond with steep reductions in the probability of voting yes on consumer protection-oriented AI legislation (the reduction hovers near a magnitude of 20 percent across the entire range of lawmaker ideology). Switching to the figure displaying partisan influence, Democratic lawmakers are much more likely to vote yes on consumer protection-oriented AI legislation than are their Republican colleagues; moreover, the tightness of the predicted probability estimate for Democrats suggests a high degree of orthodoxy in the Democratic Party regarding support for consumer protection-oriented AI legislation. While greater dispersion around the Republican estimate suggests that some Republicans support consumer protection-oriented AI legislation, the reduction in yes votes in the Republican Party compared to the Democratic suggests that the parties do view support for consumer protection-oriented AI legislation differently from one another. In short, differences in support for consumer protection-oriented AI policy appear to exist based on ideology and party.

Conclusions and next steps

In this paper, we have sought to shed light on the political factors associated with legislative activity on the fledgling AI regulation issue area. AI is itself an expansive concept, with effects spanning from research and development to civil rights and privacy to employment and competition and beyond. Yet in spite of the breadth of AI, we find a handful of political and economic factors to be consistently associated with AI policymaking and legislative support. First, economic factors like unemployment and inflation are important to consider, as we find that unemployment and inflation are negatively associated with adoption of AI policies. Future research should further investigate the mechanism

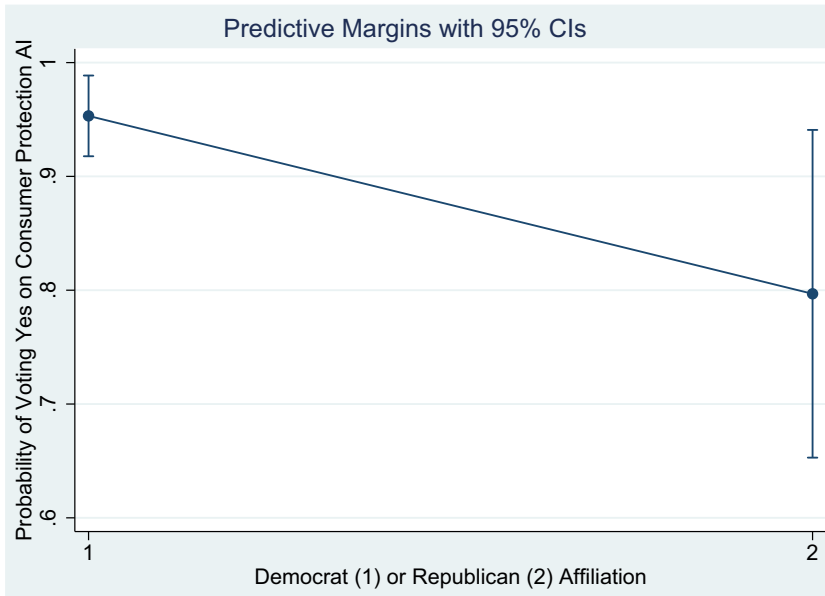


Figure 6. Influence of partisan affiliation on consumer protection AI voting.

behind this association, and it seems plausible that lawmakers today remain worried about the optics of embracing AI policy given economic distress. There is some support for the idea that unified Democratic government predicts AI policy adoption. At the individual-legislator level, ideological and partisan factors predict AI policy support, with liberal legislators supporting consumer protection AI policy and Democratic legislators supporting the same. A clear ideological and partisan consensus appears not to exist with respect to non-consumer protection-oriented AI policy.

Our research underscores the political importance of the distinction between consumer protection AI legislation and business- or economy-focused legislation. As with other sorts of business regulations, progressives and Democrats are pre-disposed to support legislation that protects the civil rights and liberties of consumers. However, at least in the abstract, this need not be the case: indeed, small-government conservatives and libertarians—as well as those concerned with various sorts of “algorithmic bias”—may have reason to support certain sorts of consumer protection legislation. With a handful of exceptions (Idaho being the most notable), Republican states are not especially likely to pass consumer protection legislation, and it remains to be seen how Republican states will approach the complex issue of how to regulate AI as the development and application of AI technology moves forward.

Our hope is that future research will continue to track these policy differences, and the partisan and ideological dynamics underlying their adoption, as AI policy matures as an issue area. Given the inherent tensions of AI policy, including privacy, economic growth, security, and innovation, AI’s precise partisan and political cleavages are bound to evolve over time. As they do, the sorts of policy changes and solutions that are possible in US legislatures—and the regulations to which businesses domestic and international are subject—are likely to shift accordingly.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/bap.2023.40>.

References

- Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.
- Acemoglu, Daron, and Pascual Restrepo. 2020. “The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand.” *Cambridge Journal of Regions, Economy and Society* 13 (1): 25–35.
- Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2017. “Revisiting the Risk of Automation.” *Economics Letters* 159: 157–160.

- Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29 (3): 3–30.
- Boehmke, Frederick J., and Paul Skinner. 2012. "State Policy Innovativeness Revisited." *State Politics & Policy Quarterly* 12 (3): 303–329.
- Box-Steffensmeier, Janet M., and Christopher Zorn. 2002. "Duration Models for Repeated Events." *The Journal of Politics* 64 (4): 1069–1094.
- Carter, Denise. 2020. "Regulation and Ethics in Artificial Intelligence and Machine Learning Technologies: Where Are We Now? Who Is Responsible? Can the Information Professional Play a Role?" *Business Information Review* 37 (2): 60–68.
- Chae, Yoon. 2020. "US AI Regulation Guide: Legislative Overview and Practical Considerations." *The Journal of Robotics, Artificial Intelligence & Law* 3 (1): 17–40.
- de Almeida, Patricia Gomes Rêgo, Carlos Denner dos Santos, and Josivania Silva Farias. 2021. "Artificial Intelligence Regulation: A Framework for Governance." *Ethics and Information Technology* 23 (3): 505–525.
- Felz, Daniel, and Kimberly Peretti. 2022. AI Regulation in the U.S.: What's Coming and What Companies Need to Do in 2023. *Newsletter*. <https://www.alston.com/en/insights/publications/2023/02/ai-regulation-in-the-us>.
- Fosch-Villaronga, Eduard, and Michiel Heldeweg. 2018. "Regulation, I Presume?' Said the Robot—Towards an Iterative Regulatory Process for Robot Governance." *Computer Law & Security Review* 34 (6): 1258–1277.
- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114: 254–280.
- Garlick, Alex. 2023. "Laboratories of Politics: There Is Bottom-up Diffusion of Policy Attention in the American Federal System." *Political Research Quarterly* 76 (1): 29–43.
- Grumbach, Jacob. 2022. *Laboratories Against Democracy: How National Parties Transformed State Politics*, vol. 184. Princeton, NJ: Princeton University Press.
- Gurkaynak, Gonenc, Ilay Yilmaz, and Gunes Haksever. 2016. "Stifling Artificial Intelligence: Human Perils." *Computer Law & Security Review* 32 (5): 749–758.
- Holder, Chris, Vikram Khurana, Faye Harrison, and Louisa Jacobs. 2016. "Robotics and Law: Key Legal and Regulatory Implications of the Robotics Age (Part I of II)." *Computer Law & Security Review* 32 (3): 383–402.
- Jones, Bradford S., and Regina P. Branton. 2005. "Beyond Logit and Probit: Cox Duration Models of Single, Repeating, and Competing Events for State Policy Adoption." *State Politics & Policy Quarterly* 5 (4): 420–443.
- NCSL. 2022. Report: Legislation Related to Artificial Intelligence. National Conference of State Legislatures. <https://www.ncsl.org/technology-and-communication/legislation-related-to-artificial-intelligence>.
- New York Times Editorial Board. 2012. "The Bills to Nowhere." *The New York Times*. <https://www.nytimes.com/2012/06/08/opinion/the-bills-to-nowhere.html>.
- O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- OxfordInsights. 2022. Government AI readiness index. <https://oxfordinsights.com/ai-readiness/ai-readiness-index/>.
- Pack, Harper. 2022. "Regulation of Artificial Intelligence in the United States." *Interactive Robotics: Legal, Ethical, Social and Economic Aspects: Selected Contributions to the INBOTS Conference 2021, 18–20 May, 2021*. Springer, 233–237.
- Parinandi, Srinivas. 2023. *Following in Footsteps or Marching Alone? How Institutional Differences Influence Renewable Energy Policy*. Ann Arbor, MI: University of Michigan Press.
- Poole, Keith T., and Howard Rosenthal. 1997. *Congress: A Political-Economic History of Roll Call Voting*. New York, NY: Oxford University Press.
- Reed, Chris. 2018. "How Should We Regulate Artificial Intelligence?" *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376 (2128): 20170360.
- Rice, Stuart A. 1927. "The Identification of Blocs in Small Political Bodies." *American Political Science Review* 21 (3): 619–627.
- Shor, Boris, and Nolan McCarty. 2011. "The Ideological Mapping of American Legislatures." *American Political Science Review* 105 (3): 530–551.
- Sorkin, Aaron, Ravi Mattu, Bernhard Warner, Michael de la Merced, Lauren Hirsch, and Ephrat Livni. 2023. Why Lawmakers Aren't Rushing to Police A.I. *New York Times* (3 March 2023).
- Squire, Peverill. 2017. "A Squire Index Update." *State Politics & Policy Quarterly* 17 (4): 361–371.
- The White House. 2023. Blueprint for An AI Bill of Rights. Office of Science, Technology, and Policy. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.
- Tutt, Andrew. 2017. "An FDA for Algorithms." *Administrative Law Review* 69: 83.
- Wright, Scott A., and Ainslie E. Schultz. 2018. "The Rising Tide of Artificial Intelligence and Business Automation: Developing an Ethical Framework." *Business Horizons* 61 (6): 823–832.
- Zardoya Jiménez, Natalia, and Christi Amesti Mendizábal. 2022. "Can the Legal Requirements Foreseen for Other Sectors Aimed at Protecting the Market and, Ultimately, the People and Users Themselves, Be Transferred to the Robotics Industry Sector?" *Interactive Robotics: Legal, Ethical, Social and Economic Aspects: Selected Contributions to the INBOTS Conference 2021, 18–20 May, 2021*, Springer, 199–206.