

The Bottleneck in AI Governance and Oversight: Evidence from 1,419 State Bills

Anonymous Authors¹

Abstract

Understanding what makes AI governance succeed or stall has immediate practical consequences for society. Since 2017, U.S. state AI bill volume has grown more than 35-fold, yet most never reach a floor vote. We present the first state-level bill outcome analysis to decompose the AI governance bottleneck into two distinct gates, analyzing 1,419 bills across all 50 states from 2017 to 2026. Our two-stage hurdle model reveals the bottleneck is not monolithic: specialist committee referral is the primary barrier to a floor vote, while cross-party alignment is the decisive factor at enactment. States are enacting AI laws at 8× the federal rate, making them the viable near-term venue for AI governance. For the ML and AI community, our findings translate directly into where and how the AI community should engage in the legislative processes to advance trustworthy AI governance.

1. Introduction

AI governance and oversight require legislation that clears committee, reaches a floor vote, and gets enacted. At the federal level, U.S. Congressional AI bills from 2017 to 2025 showed only 2.3% were enacted, with 59.3% stalling in committees before any floor vote (Khan et al., 2025). Whether this bottleneck is specific to federal institutions or reflects broader structural properties of U.S. legislative processes is an open question. If the bottleneck persists across state legislatures, the obstacle to AI governance runs deeper than any single institutional design.

Existing work on state-level AI legislation is limited. Parinandi et al. (2024) examine roll-call voting patterns across all 50 states while Dawson et al. (2026) apply comparative

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

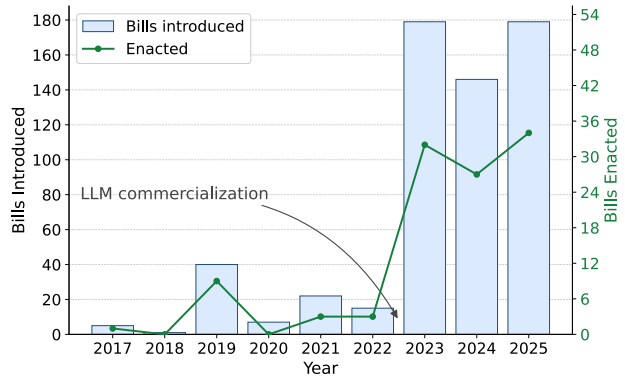


Figure 1. U.S. state legislatures, AI bill volume (bars, left axis) and enacted count (line, right axis) in completed sessions, 2017-2025. The 2023 volume surge coincides with the commercialization of large language models.

analysis to bill types. However, no prior work has conducted bill-level outcome modeling at scale across all 50 states over the last nine-year period.

Recent events illustrate the AI governance and policy stakes. Utah’s Mental Health Chatbot Act (HB 452), enacted March 2025, established the first state regulatory framework for GenAI mental health tools. Months later, a teen suicide linked to an AI companion led to the lawsuit *Raine v. OpenAI* (CNN, 2025), generating constituent pressure across multiple states: California alone enacted three companionbot and AI safety bills, with similar legislation in session in Washington, New York, and Oregon (Stanford Human-Centered AI, 2026; Washington, 2024; California, 2024). Understanding what makes such legislation succeed or stall has immediate practical consequences.

In January 2025, U.S. President’s Executive Order 14365 attempted to prohibit state AI regulation (The White House, 2025); the Senate subsequently removed the proposed 10-year moratorium on state legislation (U.S. Senate, 2025), affirming states as the active regulatory layer during a period of federal gridlock. Internationally, the EU AI Act has encountered significant operationalization delays, where as of May 2025, only 3 of 27 member states had designated national competent authorities, and only one had an operational AI regulatory sandbox (AI Act, 2025).

This paper examines the bottleneck in AI governance and oversight at the enforceable layer in U.S., the state legislature, covering all 50 U.S. states (2017-2026), and makes four contributions:

- We show, state AI bill volume has grown more than 35-fold since 2017, accelerating with AI capability milestones (Figure 1), while states enact AI law at 8× the federal rate; moving substantially faster on AI governance.
- A novel two-stage hurdle model reveals the bottleneck is not monolithic: specialist committee referral is the primary barrier that hinders access to floor, while cross-party coalitions are the decisive factor for enactment. This shows, the two gates are structurally different problems requiring different strategies.
- We translate findings into concrete action for the AI community, deriving stage-specific recommendations to engage specialist committees through testimonials, build cross-party co-sponsorship at the drafting stage, and prefer generalist committee pathways to directly help AI governance succeed.
- We release the first comprehensive cross-state labeled legislative dataset of 50 U.S. states publicly.

2. Related Work

AI legislative analysis. Parinandi et al. (2024) examine the politics and content of state AI legislation using regression on roll-call votes across all 50 states, finding that Democratic-controlled governments predict AI policy adoption and that partisan divergence differs between consumer protection and economic development bills. More recently, Dawson et al. (2026) applied qualitative comparative analysis (QCA) to 386 bills, finding responsible governance bills passed at 38.6% vs. 15.5% for transparency bills. However, none models bill progression through legislative gates. Our work extends this line with a two-stage hurdle model decomposing structurally distinct gates and identifying features that correlate to legislative progress at each.

The Stanford HAI AI Index 2026 (Stanford Human-Centered AI, 2026) provides the most comprehensive state-level characterization of AI activity to date, documenting three dimensions of variation across US states: private investment concentration (California alone accounts for 75% of national private AI investment); AI workforce density (California, Texas, and New York together account for roughly a third of all US AI job postings); and recent legislative activity documenting state-level regulatory events and cross-state responses. While the HAI report characterizes the economic and legislative landscape across states, it does not model, cross-correlations, bill-level outcomes or

identify structural features that predict legislative success. We use its economic indicators as covariates in a two-stage hurdle model providing the first quantitative account of why AI legislation succeeds or stalls at each legislative gate.

Several legislative tracking tools, including NCSL (NCSL, 2024), MultiState (MultiState, 2026), and the Brown University AISLE Portal, provide searchable registries of state AI bills but do not compute outcome rates or predictive models.

A recent analysis of 150 federal U.S. Congressional *AI bills* from 2017-2025 showed an enactment rate of 2.3% and a stall rate of 59.3% (Khan et al., 2025). Once bills gained committee attention, structural factors (more co-sponsors correlated with more stalling) were discovered.

Legislative bottleneck theory and policy governance.

Committees serve an informational function in legislative organization (Krehbiel, 1991), but majority-party agenda control can suppress bills before they reach a floor vote, a phenomenon known as legislative pigeonholing (Cox & McCubbins, 2005). Furthermore, bills addressing technically complex or politically uncertain domains are particularly vulnerable to be buried in committees. Finally, policy adaptation across states happens in waves, with policy leaders pioneering frameworks that others adopt with modifications (Shipan & Volden, 2008) and (Berry & Berry, 1990).

3. Methods

3.1. Dataset Curation

3.1.1. DATA SOURCE

We collected all AI-related state legislation introduced between 2017 and 2026 across all 50 U.S. states, including bill metadata, sponsor information, committee assignments, vote records, and status history. We used the LegiScan API (LegiScan, 2024) querying for “artificial intelligence” as the search string, returning bills with AI as a primary or significant secondary subject in the bill meta-data.

3.1.2. SUPPLEMENTARY DATASETS

We supplement the LegiScan bill data with three state-level datasets. Legislative professionalism is measured using the Squire Index (Squire, 2023), a standardized composite of legislative salary, session length, and staff size benchmarked against Congress. AI economic context is drawn from the Stanford HAI AI Index 2026 (Stanford Human-Centered AI, 2026), providing state-level AI workforce density (share of jobs in AI occupations, 2025), cumulative public AI research grants (2013-2024), and public AI procurement via contracts and OTAs (2013-2024). Legislative majority party (2017-2026) and governor party affiliation are drawn from the National Conference of State Legislatures (NCSL,

2025) and the National Governors Association (NGA, 2025) respectively, used to construct the majority-party sponsor and bipartisan co-sponsorship features in the hurdle model.

3.1.3. OUTCOME CLASSIFICATION

We classify each bill into an outcome state (in-progress or terminal) as follows:

- **Introduced:** all bills at entry.
- **No Action:** introduced or prefiled with no activity.
- **In Committee:** referred to committee, awaiting action.
- **Calendar:** scheduled for floor consideration.
- **Advancing:** passed one or both chambers.
- **Floor Vote:** reached a floor vote.
- **Stalled:** session closed without a floor vote (terminal).
- **Failed:** floor vote did not pass (terminal).
- **Vetoed:** vetoed by governor (terminal).
- **Enacted:** signed into law (terminal).

3.1.4. SCOPE AND SAMPLES

To study specific scopes, the dataset is used to create four sample sets, each mapped to a distinct use. Table 1 summarises the derivation.

Table 1. Scope and Samples

Scope	N	Used for
All bills	1,419	Bill flow; volume trends (§4.1)
Completed sessions	594	Outcome rates (§4.2)
State comparisons	438 (14 states)	State-level rates (§4.3)
Hurdle model	541	Progression features (§4.4)

All bills ($n = 1,419$): The full dataset spans 2017 to 2026 across all 50 states. It is used for volume trend analysis and the Sankey flow diagram.

Completed sessions ($n = 594$): Bills whose legislative session has formally closed. All outcome rate metrics (enactment, stall, action) are computed on this sample only. The 825 bills in active sessions are excluded here to avoid selection bias from partially-observed sessions.

State comparisons ($n = 438$, 14 states): For an objective state-level comparison, we limit to states with ≥ 10 completed-session bills related to AI. This excludes 6 states with no completed-session bills related to AI as well as low-volume states such as AK ($n = 4$), and DE ($n = 4$), where a single AI related bill enactment would produce an apparently high enactment rate.

Hurdle model ($n = 541$): Among all the 594 completed-session bills, 53 bills do not have majority party affiliation: 5 Nebraska bills (non-partisan unicameral legislature, no majority party) and 48 bills where the primary sponsor is a committee. These were removed from this analysis.

3.2. Legislative Progress Metrics

Let N denote the total number of bills in completed legislative sessions. We define three primary outcome rates:

$$\text{Enactment rate} = \frac{|\text{Enacted}|}{N} \quad (1)$$

$$\text{Stall rate} = \frac{|\text{Stalled}|}{N} \quad (2)$$

$$\text{Action rate} = \frac{|\text{Enacted}| + |\text{Failed}| + |\text{Vetoed}|}{N} \quad (3)$$

The action rate (Equation 3) measures legislative engagement regardless of direction: a bill contributes whether it passes, fails, or is vetoed, but not if it stalls without any floor disposition. Traditionally, legislatures are known to enact only a fraction of introduced bills (Squire & Hamm, 2005).

3.3. Two-Stage Binary Hurdle Model

We implement a two-stage binary hurdle model that decomposes legislative progress into two sequential gates, each modeled as a separate logistic regression:

- **Stage 1 (Floor Gate):** Among the 541 filtered completed-session bills, predicts if a bill reaches floor vote (enacted, failed, or vetoed) or stalls across any chamber. Classes: stalled = 370, reached floor = 171.
- **Stage 2 (Enactment Gate):** Among the 171 bills that reached a floor vote, predicts enactment vs. failure or veto. Classes: failed/vetoed = 73, enacted = 98.

To identify which coalition, procedural, and institutional factors predict progression through each gate, the hurdle model includes features of three types.

Sponsor and coalition features capture the political support behind a bill: sponsor count (z-scored), Republican sponsorship (vs. Democrat reference), bipartisan status, and whether the primary sponsor belongs to the chamber majority party (*sponsor.in.majority*, coded using NCSL partisan control data for all 50 states 2017 to 2026 (NCSL, 2025)).

Procedural features capture the structural pathway: chamber of origin (Senate vs. House; Assembly chambers mapped to House) and bill type (Resolution vs. Bill).

Institutional context features capture the broader political environment: governor’s party alignment with the chamber majority (governor_aligned, from NGA official records for all 50 states 2017 to 2026 (NGA, 2025)), and seven committee-category indicators (Tech/AI; Judiciary/Consumer; Finance/Appropriations; Rules/Procedure; Commerce/Economic; Education/Health/Labor; Government Operations; reference: Other/Unclassified), which

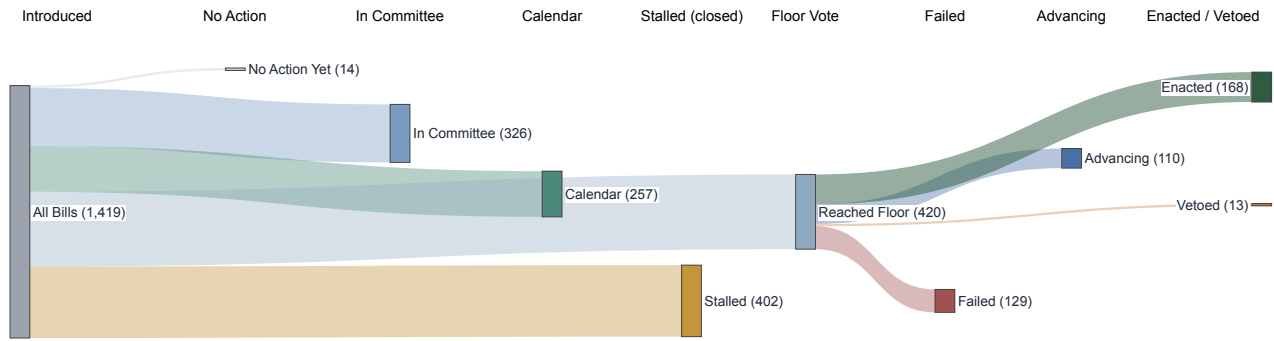


Figure 2. Bill flow diagram showing the progression of 1,419 AI-related state bills (2017 to 2026) from introduction through final outcomes.

proxy for the policy domain and the reviewing committee’s likely expertise and receptiveness.

Statistical estimation: Point estimates are full-sample MLEs: each logistic regression is fit on the complete filtered sample ($n = 541$ for Stage 1, $n = 171$ for Stage 2). To obtain confidence intervals and p-values without assuming asymptotic normality, we use 2,000 bootstrap iterations with replacement. Each iteration draws a bootstrap sample of size n (with replacement), fits the model, and evaluates on the out-of-bag (OOB) observations (approximately 37% of the sample not drawn in any given iteration). Bootstrap p-values are computed as twice the proportion of bootstrap coefficients crossing zero (two-sided). Confidence intervals are the 2.5th to 97.5th percentile of the bootstrap distribution. AUC is reported as mean OOB AUC across 2,000 iterations, a cross-validated estimate that avoids overfitting bias from a single held-out split. P-values are adjusted for multiple comparisons within each stage separately using the Benjamini-Hochberg (BH) False Discovery Rate procedure (14 predictors per stage; FDR controlled at $q < 0.05$ per stage independently). We use L2 (ridge) regularization ($C = 1.0$), and balance the class weights to prevent trivial majority-class prediction (Stage 1: 370 stalled vs. 171 floor-reached; Stage 2: 73 failed/vetoed vs. 98 enacted). Data and source code will be made publicly available at www.anonymized.com.

4. Results

4.1. Bill Flow Overview

Figure 2 displays the full flow of all 1,419 bills across the ten outcome categories defined in Section 3.1.3. Of these, 168 have already been enacted. This count is expected to grow as 825 bills are in active sessions, including 257 on the legislative calendar, 326 in committee, and 14 with no

action yet. Failure is distributed across multiple stages: 402 bills stalled before a floor vote, 129 failed at vote, and 13 were vetoed, indicating that attrition is not confined to a single gate but accumulates across the pipeline. In order to isolate enactment and failure rates from the impact of in-progress bills, we examine the 594 completed-session bills in detail next.

4.2. Completed-Session Outcomes

We next focus on the 594 bills in completed sessions to isolate definitive outcomes, from the bills that are still in progress. Among these 594, 18.4% were enacted while remaining didn’t due to stalling (67.7%), failed vote (12.3%) or being vetoed (1.7%), see Table 2. This data identifies two major failure modes: First, 402 bills stalling before floor vote, and second 73 failing at vote, indicating that the legislative bottleneck operates at two structurally different stages. We study this bi-fractional split in Section 4.4.

Table 2. Outcomes in Completed Sessions ($n = 594$)

Outcome	Count	%
Enacted	109	18.4
Stalled (no floor vote)	402	67.7
Failed (floor vote)	73	12.3
Vetoed	10	1.7
Total	594	100.0

Action Rate: The state legislatures action rate on AI bills (those receiving any definitive legislative disposition) was 32.3%, 7.6× higher than 4.23% (Khan et al., 2025) at the federal level.

Enactment Rate: Correspondingly, state legislatures also achieved a *substantially higher* enactment rate on AI related bills (18.4% vs. 2.3%) than Congress, indicating that state procedural structures are less restrictive than the federal

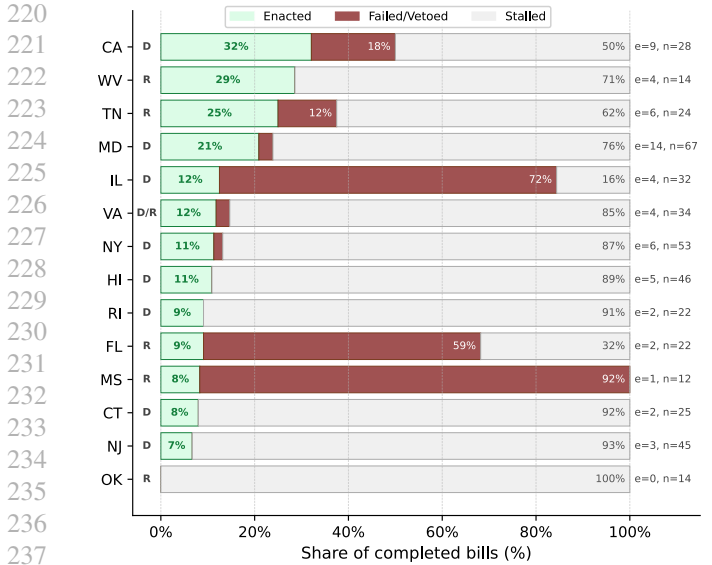


Figure 3. State AI bill enactment rates for the 14 states with ≥ 10 completed-session bills. (e) indicates enacted and (n) completed-session bills. D/R represents majority party during most of the corresponding legislative period in the state.

level for AI legislation. The analysis shows that U.S. States had a $8\times$ higher enactment rate than Congress for AI related bills during 2017-2026. In comparison, Quorum (2026) reports that U.S. states enactment rate on all bills in the last decade was $4.7\times$ higher than Congress (25.84% State enact. rate vs. 5.5% Congress enact. rate).

4.3. State-Level Variation

Figure 3 reports completed-session rates for the 14 states with ≥ 10 completed-session bills. The seven states with the most completed-session bills are all Democratic or split-controlled. We found, that legislative professionalism (Squire Index, Squire (2023)) shows no significant correlation with enactment rate across the 14-states ($r_s = -0.03$); and examine the heterogeneous patterns next.

4.3.1. HETEROGENEOUS LEGISLATIVE MECHANISMS

Top Enactors: California leads all states in enactment rate (32.1%, Figure 3), with 6 of its 9 enacted bills constituting substantive regulation: training data transparency (AB2013), generative AI accountability (SB896), AI disclosure in political advertisements (AB2355), and AI in healthcare decisions (AB3030), among others. West Virginia (28.6%) and Tennessee (25.0%) rank second and third, but their enacted bills mostly cover study commissions, task forces and AI-generated sexual content prohibitions, and not regulations.

Top Floor Failed: Illinois (72% failed), Mississippi (92% failed), and Florida (59% failed) are states where bills fail




Figure 4. AI workforce density (share of state jobs in AI, 2025; (Stanford Human-Centered AI, 2026)) versus enactment rate for 14 states with ≥ 10 completed-session bills. Three states (WV, TN, VA) are excluded from the correlation on substantive grounds (see text); the linear trend and Spearman r_s are shown for 11 states.

at the floor vote stage. Illinois bills that reached the floor covered substantive and politically topics that could result in significant unresolved debates, including labor and employment AI (algorithmic hiring decisions) and consumer protection (AI labeling, insurance oversight) (see Appendix Table 4). On the other hand, the failed bills in MS and FL span both regulatory and non-regulatory proposals suggesting a general reluctance towards AI governance.

Top Committee Stall: Connecticut (92% stall), New Jersey (93.3% stall), and Oklahoma (100% stall) all show extreme committee stall rates. In CT and NJ, every bill that cleared committee was enacted (100% floor pass rate), likely because leadership pre-cleared bills before they reach the floor (Cox & McCubbins, 2005; Sinclair, 2016). Oklahoma’s all stalled bills span diverse topics including synthetic media and AI governance (Table 5), likely reflecting diffusion of model-legislation templates from other states. However, the absence of a commercial AI ecosystem generates little local pressure for enactment: the HAI AI Index 2026 reports no mapped private AI investment in Oklahoma (Stanford Human-Centered AI, 2026).

4.3.2. ECONOMIC PROFILE VS. LEGISLATIVE ACTION

Figure 4 shows that AI workforce density, a proxy of economic context, correlates ($r_s = 0.70$, $p = 0.016$) with state enactment rate. AI workforce density captures the share of state jobs in AI occupations in 2025 (Stanford Human-Centered AI, 2026), representing the working population engaged in building AI systems that are the object of regulation. States with a higher share of workers in AI occupations

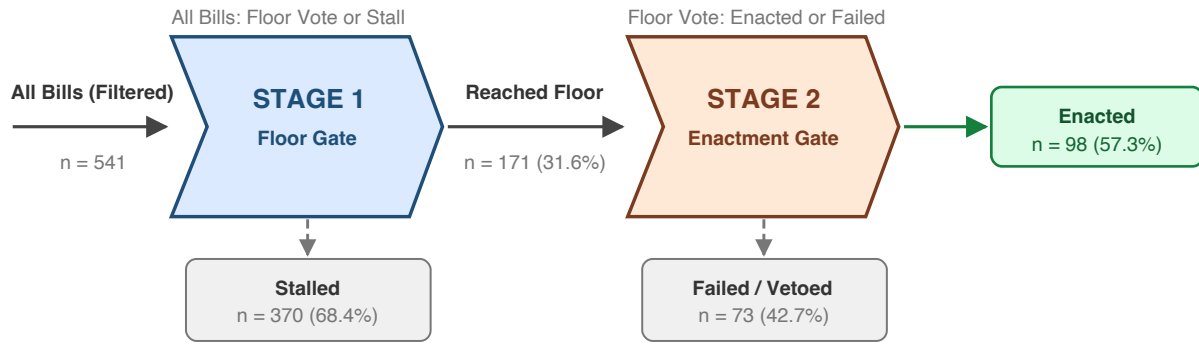


Figure 5. Sequential binary hurdle model of AI bill legislative progress, with two gates: floor vote (Stage 1) and enactment (Stage 2). Bills must clear both gates to become law.

produce more AI output, which in turn generates both the expertise needed to draft technically credible legislation and the constituent pressure to enact it. This result is exploratory given $n = 11$, indicating that states where AI development is most concentrated show greater legislative progress.

We excluded the three states from the correlation analysis, that show patterns plausibly driven by factors outside commercial AI ecosystem. West Virginia (28.6% enactment, $n = 14$) and Tennessee (25.0%) reflect Republican trifecta legislatures focused on AI promotion rather than regulation; Virginia (11.8%) has high AI workforce density driven largely by federal defense procurement (\$1.09B in contracts and OTAs, 2013-2024, representing nearly 60% of total US contract and OTA AI spending alongside California and Maryland (Stanford Human-Centered AI, 2026)) rather than commercial AI industry concentration, with the Department of Defense accounting for 74.1% of all federal AI contract spending. We found that investment and research spending, which fund upstream inputs rather than deployed systems, show weaker correlations with enactment rate (Appendix E).

4.4. Two-Stage Hurdle Model

The two-stage hurdle model, shown in Figure 5, decomposes legislative progress into two sequential binary decisions, each capturing a structurally distinct gate. Stage 1 (Floor Gate, $n = 541$) asks whether a bill reaches floor vote, separating the 171 bills that did from the 370 that stalled. For bills that reach the floor vote, Stage 2 (Enactment Gate, $n = 171$) asks whether a bill is enacted or not, separating the 98 enacted bills from the 73 that failed or were vetoed.

Table 3 presents coefficients and bootstrap confidence intervals for both stages, showing that the two stages capture qualitatively different political dynamics. Stage 1 shows a committee and coalition phenomenon, while Stage 2 is driven by cross-party negotiation as detailed below.

Stage 1 (Floor Gate): Features associated with reaching a floor vote. Republican sponsorship shows the strongest association ($q < 0.001$): Republican-sponsored bills are 2.19 times more likely to reach a floor vote than Democrat-sponsored bills. This is not a proxy for majority-party status; the two features are weakly and negatively correlated ($r = -0.26$), and Democrat-sponsored bills are actually more likely to have majority-party sponsorship in this sample (86.9% vs. 64.5% for Republican bills). We hypothesize that the Republican floor-access advantage may reflect bill content rather than structural position: Republican AI bills in this period may have tended toward industry-promotion and regulatory-light framings, a question that requires future study. Furthermore, sponsor count is also positively associated with floor access ($OR = 1.40$, $q = 0.014$), as broader coalitions make committee burial politically costly.

On the other hand, bills referred to specialist committees show strong negative associations with floor access: Tech/AI ($OR = 0.34$, $q < 0.001$), Judiciary/Consumer ($OR = 0.44$, $q = 0.014$), Commerce/Economic ($OR = 0.41$, $q = 0.014$), and Education/Health/Labor ($OR = 0.44$, $q = 0.014$), all relative to Other/Unclassified committees. Two interpretations are possible and cannot be distinguished from observational data alone: 1) specialist committees may act as active gatekeepers blocking bills based on domain-specific views; or, 2) contested AI bills may naturally reach into specialist committees, making committee assignment a marker of bill complexity rather than a cause of suppression. Notably, specialist committees that do pass bills to the floor (19-26% floor-reach rate) show Stage 2 enactment rates comparable to less selective committees (58-63%), corroborating the committee bottleneck at the referral stage (Figure 7).

Stage 2 (Enactment Gate): Features associated with enactment among floor-vote bills. At the floor vote, bipartisan co-sponsorship, instead of Republican origin, is the strongest feature ($OR = 5.33$, $q < 0.001$), reflecting that

The Bottleneck in AI Governance and Oversight

Table 3. Feature coefficients, odds ratio (OR), 95% CI, and q_{BH} : Benjamini-Hochberg FDR q -values discovered from the two-stage hurdle model.

Feature	Stage 1: Floor Gate (n = 541)				Stage 2: Enactment Gate (n = 171)			
	Coef.	OR	95% CI _{coef}	q_{BH}	Coef.	OR	95% CI _{coef}	q_{BH}
Sponsor & coalition features								
Sponsor count (SD)	+0.338	1.40	[+0.115, +0.607]	*0.014	+0.229	1.26	[-0.192, +1.050]	0.649
Sponsor in majority party	+0.214	1.24	[-0.252, +0.759]	0.383	+1.206	3.34	[+0.525, +1.835]	**0.005
Bipartisan	+0.396	1.49	[-0.167, +0.934]	0.258	+1.673	5.33	[+0.962, +2.481]	***<0.001
Republican (vs. Democrat)	+0.785	2.19	[+0.368, +1.248]	***<0.001	+0.079	1.08	[-0.534, +0.741]	0.930
Procedural features								
Chamber: Senate (vs. House)	+0.387	1.47	[+0.003, +0.785]	0.100	+0.037	1.04	[-0.536, +0.635]	0.930
Resolution (vs. Bill)	+0.385	1.47	[-0.230, +1.001]	0.293	+2.347	10.45	[+1.822, +2.786]	***<0.001
Institutional context features								
Governor aligned w/ majority	+0.317	1.37	[-0.213, +0.982]	0.301	+0.023	1.02	[-0.912, +0.834]	0.930
Committee: Tech/AI	-1.092	0.34	[-1.693, -0.548]	***<0.001	-0.642	0.53	[-1.416, +0.208]	0.339
Committee: Judiciary/Consumer	-0.826	0.44	[-1.450, -0.248]	*0.014	+0.070	1.07	[-0.756, +0.833]	0.930
Committee: Finance/Approp	-0.351	0.70	[-0.967, +0.239]	0.293	+0.141	1.15	[-0.646, +0.966]	0.930
Committee: Rules/Procedure	+0.248	1.28	[-0.289, +0.855]	0.383	-1.202	0.30	[-1.888, -0.537]	*0.011
Committee: Commerce/Economic	-0.902	0.41	[-1.628, -0.285]	*0.014	-0.063	0.94	[-0.907, +0.875]	0.930
Committee: Edu/Health/Labor	-0.814	0.44	[-1.470, -0.216]	*0.014	-0.179	0.84	[-0.952, +0.743]	0.930
Committee: Government Ops	-0.590	0.55	[-1.351, +0.097]	0.159	-0.272	0.76	[-1.068, +0.634]	0.930
n	541				171			
Class 0 / Class 1	370 / 171				73 / 98			

Reference categories: Bill type=Bill; Committee=Other/Unclassified; Chamber=House; Party=Democrat. Full bootstrap p-values are given in Table 9. * $q < 0.05$, ** $q < 0.01$, *** $q < 0.001$.

bills with cross-party support are 5.33 times more likely to be enacted. Examining the raw counts directly, bipartisan bills converted 85% of floor votes into law (47/55), against 45% for non-bipartisan bills (55/121), corroborating the model association (see Figure 6). Analogously, majority-party sponsorship also shows a strong positive association with enactment (OR = 3.34, $q = 0.005$). Together, these suggest that floor success may reflect broadly negotiated cross-party consensus rather than partisan momentum alone.

Bills referred to Rules or Procedure committees show the inverse pattern, with high floor-reach but substantially lower

enactment odds (OR = 0.30, $q = 0.011$). Rules committees function as majority-party agenda-control instruments, regulating floor access through special procedural rules rather than substantive deliberation (Cox & McCubbins, 2005). Bills routed this way bypass the process where cross-chamber consensus is built, arriving at the floor vote without the coalition alignment that predicts enactment. Sinclair (2016) documents that such pathways have become common for contested legislation, substituting leadership brokerage for traditional committee consensus. Finally, resolutions enacting at higher rates (OR = 10.45, $q < 0.001$) is a procedural artifact due to their non-binding nature rather than a substantive AI-policy finding.

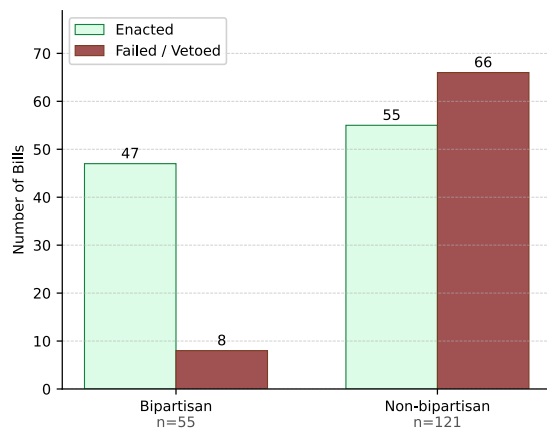


Figure 6. Floor vote outcomes by bipartisan status (Stage 2, $n = 171$). Enacted bills: Bipartisan 85%; non-bipartisan: 45%.

5. Discussion

5.1. Interpreting the Bottleneck

Two structurally distinct gates require different strategies. The two-stage hurdle model reveals that the AI governance bottleneck is not a single obstacle but two structurally distinct gates. Floor-vote access is a committee and coalition phenomenon; enactment is a cross-party consensus phenomenon. These gates differ in kind, as a strategy effective at Stage 1 (broad co-sponsorship, generalist committee referral) is structurally different from what Stage 2 requires (bipartisan framing, majority-party primary sponsor).

Why state legislatures. The $8\times$ higher state action rate reflects structural features of state legislatures: shorter and

The Bottleneck in AI Governance and Oversight

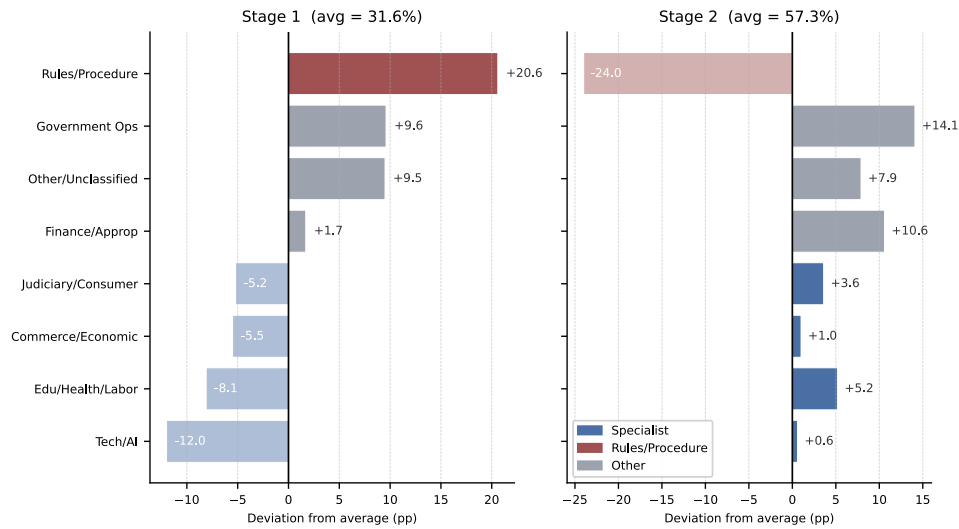


Figure 7. Committee type deviation from average at Stage 1 (floor-reach) and Stage 2 (enactment). Specialist committees (blue) sit below average at Stage 1 but close to average at Stage 2; Rules/Procedure (red) shows the inverse pattern.

time-bounded sessions create resolution pressure, smaller chambers reduce the durability of committee pigeonholing, and a narrower legislative agenda concentrates attention on fewer competing bills. These conditions make state legislatures a more tractable near-term venue for AI governance.

5.2. AI Policy Recommendations

The two-stage structure of legislative outcomes carries concrete recommendations for the AI community, safety advocates, and state policymakers working to close the gap between AI development and governance.

Engage at the specialist committee stage. ML practitioners and AI safety organizations should direct technical testimony to committee hearings to help bills move to floor vote. Specialist committees are the primary stalling point for AI governance bills: referral to a Tech/AI committee cuts floor-vote probability by two-thirds (OR = 0.34, Stage 1).

Build bipartisan co-sponsorship. Among bills that reach a floor vote, bipartisan bills are more than five times more likely to be enacted into law (OR = 5.33, $q < 0.001$, Stage 2). For AI governance to succeed, ML practitioners should frame their testimony in non-partisan way and AI policy makers should pursue active bipartisan coalition.

Prefer generalist committee pathways for initial passage. Bills routed through Rules/Procedure committees show higher floor-reach but lower enactment rates. Commerce or general government committees offer a more navigable Stage 1 pathway for AI oversight legislation.

Consider sunset provisions to lower the political cost. Regulatory AI bills face heightened committee resistance

due to perceived scope and permanence. Sunset clauses with mandatory review periods reduce the political cost of initial approval while preserving the option for renewal, making AI governance bills more viable in legislatures.

6. Conclusion

Across 1,419 state AI bills, the AI bill volume has surged 35-fold, yet enactment rates have remained flat. The bottleneck in AI governance is structural and not a shortage of bills. The two-stage hurdle model localises it: committee gatekeeping at Stage 1 and cross-party coordination failure at Stage 2. These are distinct problems requiring distinct interventions, procedural reforms to prevent committee burial, and deliberate bipartisan coalition-building before introduction. For the ML and AI community, these findings translate directly into where and how they should engage with state legislative processes to advance trustworthy AI governance.

Limitations

- Query coverage.** Using “artificial intelligence” as the search term may miss AI-adjacent legislation using alternative terminology. Total AI-adjacent activity is likely higher than our counts suggest. However, this approach provides an objective criterion.
- Companion bills.** To retain all legislative actions, we keep companion bills as separate observations; 30 such pairs appear in completed sessions. Deduplicating all 30 would shift the enactment rate from 18.4% to 18.3% and the stall rate from 67.7% to 67.0%, neither of which materially affects any reported finding.

References

Benjamini, Y. and Hochberg, Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300, 1995. doi:10.1111/j.2517-6161.1995.tb02031.x.

Berry, F. S. and Berry, W. D. State lottery adoptions as policy innovations: An event history analysis. *American Political Science Review*, 84(2):395–415, 1990. doi:10.2307/1963526.

Cihon, P., Maas, M. M., and Kemp, L. Should artificial intelligence governance be centralised? Design lessons from history. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 2020. doi:10.1145/3375627.3375857.

Cox, G. W. and McCubbins, M. D. *Setting the Agenda: Responsible Party Government in the U.S. House of Representatives*. Cambridge University Press, 2005.

Sinclair, B. *Unorthodox Lawmaking: New Legislative Processes in the U.S. Congress*. CQ Press, 5th edition, 2016.

Anonymous. Position: Bridge the gap between AI development and regulation. *Under review at ICML 2026*, 2026.

Krehbiel, K. *Information and Legislative Organization*. University of Michigan Press, 1991. doi:10.3998/mpub.8850.

LegiScan. LegiScan legislative data API, 2024. <https://legiscan.com/legiscan>.

National Conference of State Legislatures. *Artificial Intelligence 2024 Legislation*, 2024. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation>.

OECD. *OECD AI Policy Observatory*, 2024. <https://oecd.ai/en/>.

Shipan, C. R. and Volden, C. The mechanisms of policy diffusion. *American Journal of Political Science*, 52(4):840–857, 2008. doi:10.1111/j.1540-5907.2008.00346.x.

Yano, T., Smith, N. A., and Wilkerson, J. D. Textual predictors of bill survival in congressional committees. In *Proceedings of NAACL-HLT 2012*, 2012.

Riker, W. H. *The Theory of Political Coalitions*. Yale University Press, 1962.

Squire, P. and Hamm, K. E. *101 Chambers: Congress, State Legislatures, and the Future of Legislative Studies*. Ohio State University Press, 2005.

FiscalNote. 2025 State Sessions Recap: The Busiest & Most Effective States. *FiscalNote Reports*, July 2025. <https://fiscalnote.com/reports/2025-state-sessions-recap>. Accessed: May 2, 2026.

Quorum Analytics. State Legislatures vs. Congress: Which Is More Productive? *Quorum Data-Driven Insights*, February 2026. <https://www.quorum.us/data-driven-insights/state-legislatures-versus-congress-which-is-more-productive/>. Accessed: May 2, 2026.

Squire, P. A squire index update. *State Politics & Policy Quarterly*, 17(4):361–371, 2017.

Squire, P. A Squire index update: Stability and change in legislative professionalization, 1979–2021. *State Politics & Policy Quarterly*, 2023. doi:10.1017/spq.2023.22.

National Conference of State Legislatures. *State Partisan Composition*, 2025. <https://www.ncsl.org/about-state-legislatures/state-partisan-composition>. Accessed: 2026-05-02.

Stanford Human-Centered AI. *AI Index Report 2026*, 2026. https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf.

Washington State Legislature. Washington HB 2225: Concerning artificial intelligence, 2024. <https://app.leg.wa.gov/billsummary?BillNumber=2225&Year=2023>.

California State Legislature. California SB 1047: Safe and Secure Innovation for Frontier Artificial Intelligence Models Act, 2024. https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1047.

National Governors Association. *Governors*, 2025. <https://www.nga.org/governors/>. Accessed: 2026-05-02.

Cragg, J. G. Some statistical models for limited dependent variables with application to the demand for durable goods. *Econometrica*, 39(5):829–844, 1971. doi:10.2307/1909582.

Krutz, G. S. Issues and institutions: “Winnowing” in the U.S. Congress. *American Journal of Political Science*, 49(2):313–326, 2005. doi:10.1111/j.0092-5853.2005.00125.x.

Volden, C. States as policy laboratories: Emulating success in the Children’s Health Insurance Program. *American Journal of Political Science*, 50(2):294–312, 2006. doi:10.1111/j.1540-5907.2006.00185.x.

- 495 Walker, J. L. The diffusion of innovations among the Amer-
496 ican states. *American Political Science Review*, 63(3):
497 880–899, 1969. doi:10.1017/s0003055400258644.
- 498
499 Brandeis, L. D. Dissenting opinion, *New State Ice Co.*
500 *v. Liebmann*. 285 U.S. 262, 311 (1932) (Brandeis, J.,
501 dissenting), 1932.
- 502
503 European Parliament and Council. Regulation (EU)
504 2024/1689 laying down harmonised rules on artificial
505 intelligence (Artificial Intelligence Act). *Official Journal*
506 *of the European Union*, L:1–144, 2024. Entered into
507 force 1 August 2024.
- 508
509 AI Act. National implementation plans, 2025. [https://artificialintelligenceact.eu/natio-
510 nal-implementation-plans/](https://artificialintelligenceact.eu/national-implementation-plans/).
- 511
512 The White House. Ensuring a national policy framework
513 for artificial intelligence, January 2025. [https://ww
514 w.whitehouse.gov/presidential-actions
515 /2025/01/removing-barriers-to-america
516 n-leadership-in-artificial-intellige
517 nce/](https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/). Executive Order 14365, January 20, 2025.
- 518
519 U.S. Senate Committee on Commerce, Science, and Trans-
520 portation. Senate strikes AI moratorium from budget
521 reconciliation bill in overwhelming 99-1 vote, July 2025.
522 [https://www.commerce.senate.gov/2025
523 /7/senate-strikes-ai-moratorium-fro
524 m-budget-reconciliation-bill-in-overw
525 helming-99-1-vote/](https://www.commerce.senate.gov/2025/7/senate-strikes-ai-moratorium-from-budget-reconciliation-bill-in-overwhelming-99-1-vote/).
- 526
527 World Health Organization. Guidelines for implementa-
528 tion of article 13 of the WHO FCTC, 2011. Decision
529 FCTC/COP3(12). [https://fctc.who.int/res
530 ources/publications/m/item/guideline
531 s-for-implementation-article-13](https://fctc.who.int/resources/publications/m/item/guideline-s-for-implementation-article-13).
- 532
533 UNCTAD. Data protection and privacy legislation world-
534 wide, 2023. [https://unctad.org/page/data
535 -protection-and-privacy-legislation-w
536 orldwide](https://unctad.org/page/data-protection-and-privacy-legislation-worldwide).
- 537
538 Stanford Human-Centered AI. *AI Index Report 2026*, 2026.
539 [https://hai.stanford.edu/assets/file
540 s/ai_index_report_2026.pdf](https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf).
- 541
542 Khan, M. A., Akengin, M. E., and Rushdi, A. A. From
543 proposals to enactment: The procedural bottleneck in
544 AI safety regulation. In *NeurIPS 2025 Workshop on*
545 *Regulatable Machine Learning*, 2025.
- 546
547 MultiState. Artificial intelligence legislation tracker, 2026.
548 [https://www.multistate.ai/resources/
549 ai-legislation](https://www.multistate.ai/resources/ai-legislation).
- CNN Business. Parents of 16-year-old Adam Raine sue
OpenAI over ChatGPT-linked teen suicide, August 2025.
[https://www.cnn.com/2025/08/26/tech/
openai-chatgpt-teen-suicide-lawsuit](https://www.cnn.com/2025/08/26/tech/openai-chatgpt-teen-suicide-lawsuit).
Accessed: 2026-05-09.
- Parinandi, S., Crosson, J., Peterson, K., and Nadarevic, S.
Investigating the politics and content of US state artificial
intelligence legislation. *Business and Politics*, 26(2):240–
262, 2024. doi:10.1017/bap.2024.4.
- Dawson, G. S., Denford, J. S., Desouza, K. C., and Pi-
cavet, M. B. Analyzing the passage of state-level AI bills.
Brookings Institution, March 2026. [https://www.br
ookings.edu/articles/analyzing-the-p
assage-of-state-level-ai-bills/](https://www.brookings.edu/articles/analyzing-the-passage-of-state-level-ai-bills/).

A. Illinois Floor-Failed Bills by Thematic Cluster

Table 4 classifies the 23 Illinois AI bills that reached a floor vote and failed, grouped by contested policy domain. The breadth and substantive depth of these bills supports the interpretation in Section 4.3 that Illinois’s 72% floor failure rate reflects genuine legislative contestation rather than committee gatekeeping or ideological uniformity.

Table 4. Illinois AI Bills: Floor-Failed, by Thematic Cluster

Cluster	Bills
Labor & Employment AI	SB2704, SB2855, HB0364, SB0504 Algorithmic hiring, workplace AI decisions
Algorithmic Accountability	HB5322, HB4836 Impact assessments, state agency AI systems
Consumer Protection	HB5321, HB5918, HB5649 AI labeling, health insurance AI, mental health services
Voice, Likeness & Deepfake	HB3285, SB3325, SB3705, HB5848 Consent requirements, digital likeness protection
Government & Contracts	HB5099, HB5228, HB4836, HB4705 AI in government contracts, state agency reporting
Digital Content & Elections	SB2762, SB2786, HB4763 AI in political ads, digital forgeries act

B. Oklahoma Committee-Stalled Bills by Thematic Cluster

Table 5 classifies the 14 Oklahoma AI bills that stalled in committee, grouped by policy domain. Despite spanning five distinct clusters, none survived committee, consistent with the interpretation in Section 4.3 that template diffusion from other states drives introduction without local commercial pressure for enactment.

Table 5. Oklahoma AI Bills: Committee-Stalled, by Thematic Cluster ($n = 14$)

Cluster	Bills
Synthetic Media Disclosure	SB746, SB894, HB3865
AI Governance Frameworks	HB3293, HB1916, HB1917, HB1899
Health-Insurance AI	HB3577, HB1915
Ethical AI & Rights	HB3835, HB3453
Education & Government	HB3827, HB3828

C. State Groups by Economic Profile

Table 6. State groups by economic profile (14 states, ≥10 completed bills).

Group	States	Avg Enacted	Avg Stall
Tech economy	CA, IL	22.3%	32.8%
Federal/Policy corridor	MD, VA	16.3%	80.7%
Northeast + Pacific	NY, NJ, CT, RI, HI	9.2%	90.4%
Southern + Interior	TN, FL, OK, MS, WV	14.2%	53.1%

D. State Rankings by Enactment Rate

Table 7. State Rankings by Enactment Rate (States with ≥ 10 Completed Bills)

State	Party	Completed	Enacted	Enact. Rate%	Stall%
CA	D	28	9	32.1	50.0
WV	R	14	4	28.6	71.4
TN	R	24	6	25.0	62.5
MD	D	67	14	20.9	76.1
IL	D	32	4	12.5	15.6
VA	D/R	34	4	11.8	85.3
NY	D	53	6	11.3	86.8
HI	D	46	5	10.9	89.1
FL	R	22	2	9.1	31.8
RI	D	22	2	9.1	90.9
MS	R	12	1	8.3	0.0
CT	D	25	2	8.0	92.0
NJ	D	45	3	6.7	93.3
OK	R	14	0	0.0	100.0

All rates: completed-session bills only. Party: dominant legislative majority (2017–2026); VA split.

605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659

E. Public AI Investment and Enactment Rate

Public AI research grants and federal procurement spending are upstream inputs to AI development rather than direct measures of deployed AI output. We report their correlations with state enactment rates here as robustness checks. Both are positively associated with enactment after excluding the same three outlier states (WV, TN, VA) identified in the main analysis. Public research grants show a significant association ($r_s = 0.61, p = 0.048$; Figure 8), while federal contracts and OTAs show a weaker, non-significant trend ($r_s = 0.57, p = 0.067$; Figure 9). The weaker signal for procurement spending is consistent with our exclusion of Virginia, whose high procurement total reflects defense contracting rather than commercial AI concentration.

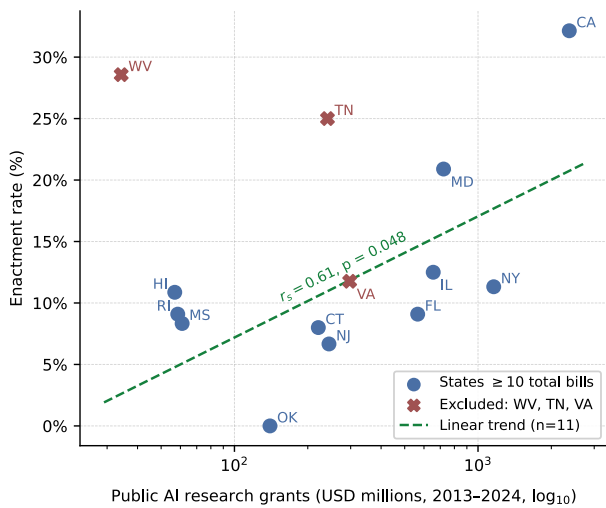
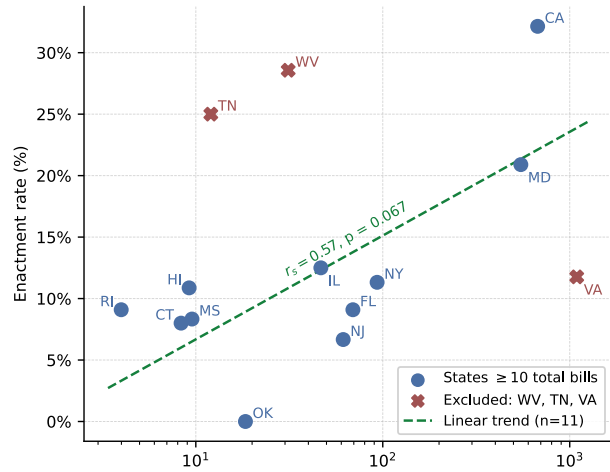


Figure 8. Public AI research grants (USD millions, 2013–2024, \log_{10} ; Stanford HAI (Stanford Human-Centered AI, 2026)) versus enactment rate for 14 states with ≥ 10 completed-session bills. Three states (WV, TN, VA) are excluded on the same substantive grounds as in the main analysis; the linear trend and Spearman r_s are shown for the remaining 11 states.



Public investment via contracts and OTAs (USD millions, 2013–2024, \log_{10})

Figure 9. Public investment in AI via government contracts and Other Transaction Agreements (OTAs, USD millions, 2013–2024, \log_{10} ; Stanford HAI (Stanford Human-Centered AI, 2026)) versus enactment rate for 14 states with ≥ 10 completed-session bills. Three states (WV, TN, VA) are excluded; the linear trend and Spearman r_s are shown for the remaining 11 states.

F. Committee Referral Patterns and Bill Outcomes

Table 8. Committee Referral Patterns and Bill Outcomes ($n = 541$ hurdle sample)

Committee	Stage 1		Stage 2	
	N	Floor%	N_{floor}	Enact.%
Tech/AI	97	19.6	19	57.9
Judiciary/Consumer	87	26.4	23	60.9
Commerce/Economic	46	26.1	12	58.3
Edu/Health/Labor	68	23.5	16	62.5
Finance/Approp	84	33.3	28	67.9
Rules/Procedure	69	52.2	36	33.3
Government Ops	34	41.2	14	71.4
Other/Unclassified	56	41.1	23	65.2
Overall	541	31.6	171	57.3

Stage 1: floor-reach rate among all 541 hurdle-sample bills.
 Stage 2: enacted rate among the 171 floor-reached bills.

G. Full Hurdle Model Results with Bootstrap P-Values

Table 9 reproduces the two-stage hurdle model with all columns, including the raw bootstrap p-values (p_{boot}) prior to BH correction.

H. AI Safety Bills: A Sub-Analysis

A central question for the AI4GOOD audience is whether AI safety legislation specifically faces distinct legislative outcomes. We identify 18 bills (3.0% of completed-session bills) whose titles or descriptions contain “safety” (case-insensitive); results are shown in Table 10.

Safety-mentioning bills enact at 38.9%, more than double the rate for other AI bills (17.7%), with a nominally significant difference ($p = 0.048$, not corrected for multiple comparisons). This analysis was exploratory and not pre-registered; the p -value is not FDR-corrected for the additional test. Stall rates do not differ significantly (61.1% vs. 67.9%, $p = 0.727$). The safety bills that *did* enact include New York budget legislation, Maryland school-safety technology provisions, and a California health-care AI bill; this heterogeneous set cautions against over-interpreting the result.

The small n ($n = 18$) limits statistical power. Notably, the sample contains no analogue to the federal finding (Anonymous, 2026) that *none* of the enacted laws focused on AI safety: at least 7 state enacted bills explicitly reference “safety,” suggesting states may be somewhat more willing to address AI safety concerns procedurally than Congress has been.

I. Political Context Variables: Sample Notes

Governor alignment. The *governor_aligned* variable is coded 1 if the governor’s party matches the chamber majority party in the year the bill was introduced, using National Governors Association official records for all 50 states 2017 to 2026. In the $n = 541$ hurdle sample, 474 bills (87.6%) were introduced in governor-aligned contexts and 67 bills (12.4%) in divided-government contexts. The low variation in this variable is consistent with the null finding at both stages ($q = 0.30$; $q = 0.93$); insufficient sample variation in the minority condition limits statistical power to detect an effect even if one exists.

Sponsor in majority party. The *sponsor_in_majority* variable is coded 1 if the primary sponsor’s party matched the controlling majority of the bill’s originating chamber (House or Senate) in the year introduced, using NCSL-compiled partisan control data. In the $n = 541$ sample, 433 bills (80.0%) were sponsored by a majority-party member and 108 bills (20.0%) by a minority-party member. Majority-party sponsorship is non-significant at Stage 1 (floor access: OR = 1.24, $q = 0.38$) but significant at Stage 2 (enactment: OR = 3.34, $q = 0.005$), indicating that institutional majority status primarily matters at the final vote rather than at the committee gatekeeping stage.

The Bottleneck in AI Governance and Oversight

Table 9. Two-Stage Hurdle Model: Full Results including Bootstrap P-Values

Feature	Stage 1: Floor Gate (n = 541)					Stage 2: Enactment Gate (n = 171)				
	Coef.	OR	95% CI _{coef}	p _{boot}	q _{BH}	Coef.	OR	95% CI _{coef}	p _{boot}	q _{BH}
Sponsor count (SD)	+0.338	1.40	[+0.115, +0.607]	0.003	0.014	+0.229	1.26	[-0.192, +1.050]	0.278	0.649
Chamber: Senate (vs. House)	+0.387	1.47	[+0.003, +0.785]	0.050	0.100	+0.037	1.04	[-0.536, +0.635]	0.909	0.930
Sponsor in majority party	+0.214	1.24	[-0.252, +0.759]	0.383	0.383	+1.206	3.34	[+0.525, +1.835]	0.001	0.005
Bipartisan	+0.396	1.49	[-0.167, +0.934]	0.166	0.258	+1.673	5.33	[+0.962, +2.481]	<0.001	<0.001
Republican (vs. Democrat)	+0.785	2.19	[+0.368, +1.248]	<0.001	<0.001	+0.079	1.08	[-0.534, +0.741]	0.724	0.930
Resolution (vs. Bill)	+0.385	1.47	[-0.230, +1.001]	0.226	0.293	+2.347	10.45	[+1.822, +2.786]	<0.001	<0.001
Committee: Tech/AI	-1.092	0.34	[-1.693, -0.548]	<0.001	<0.001	-0.642	0.53	[-1.416, +0.208]	0.121	0.339
Committee: Judiciary/Consumer	-0.826	0.44	[-1.450, -0.248]	0.005	0.014	+0.070	1.07	[-0.756, +0.833]	0.866	0.930
Committee: Finance/Approp	-0.351	0.70	[-0.967, +0.239]	0.230	0.293	+0.141	1.15	[-0.646, +0.966]	0.684	0.930
Committee: Rules/Procedure	+0.248	1.28	[-0.289, +0.855]	0.373	0.383	-1.202	0.30	[-1.888, -0.537]	0.003	0.011
Committee: Commerce/Economic	-0.902	0.41	[-1.628, -0.285]	0.004	0.014	-0.063	0.94	[-0.907, +0.875]	0.893	0.930
Committee: Edu/Health/Labor	-0.814	0.44	[-1.470, -0.216]	0.006	0.014	-0.179	0.84	[-0.952, +0.743]	0.680	0.930
Committee: Government Ops	-0.590	0.55	[-1.351, +0.097]	0.091	0.159	-0.272	0.76	[-1.068, +0.634]	0.539	0.930
Governor aligned w/ majority	+0.317	1.37	[-0.213, +0.982]	0.258	0.301	+0.023	1.02	[-0.912, +0.834]	0.930	0.930
n	541					171				
Class 0 / Class 1	370 / 171					73 / 98				
AUC-ROC (mean OOB)	0.664					0.797				

p_{boot}: proportion of 2,000 bootstrap iterations with coefficient crossing zero (two-sided). q_{BH}: Benjamini-Hochberg FDR correction within each stage (14 tests). OR = odds ratio. 95% CI on log-odds scale. Sponsor count z-scored. Reference categories: Bill type=Bill; Committee=Other/Unclassified; Chamber=House; Party=Democrat.

Table 10. AI Safety vs. Non-Safety Bills in Completed Sessions

Group	n	Enacted	Enact.%	Stalled	Stall%
Safety bills	18	7	38.9	11	61.1
Non-safety bills	576	102	17.7	391	67.9

Chi-squared (enactment): $\chi^2 = 3.91, p = 0.048$.
 Chi-squared (stall): $\chi^2 = 0.12, p = 0.727$.
 Safety group is small (n = 18); interpret with caution.