

Open-Weight AI: A Legislative Guide

→ Neil Chilson | Head of AI Policy

BOTTOM LINE

There are millions of free-to-download open weight AI models. These models let startups, nonprofits, and businesses build with capable AI without raising the capital to train a model, sharing data with third parties, or locking themselves into one vendor. Yet Congress and the states are now weighing rules that would restrict open models: preapproval, presumptions that attach to model size, and liability for what strangers do with a downloaded file. Even when AI legislation doesn't specifically target open models, compliance costs can kill them. The largest labs can pay a hefty compliance bill. Universities and startups cannot, and no model developer can control the use of open weights once released. If U.S. open models recede, foreign models will fill the gap.

CORE RULE Identify the harm, place the duty on the actor who controls it, and keep the United States the best place to build and release open models.

FOUR TERMS LAWMAKERS SHOULD KNOW

Open-weight AI: Trained model's weights (the parameters set during training) available for download. It need not disclose training data or source code.

Open-source AI: Weights plus source code and sometimes other pieces such as training data available for download.

Developer and deployer: The model developer trains and releases a model. The deployer fine-tunes it, supplies inputs and harnesses, aims it at a task, and onboards users. Laws that impose liability on developers directly impact open releases.

Openness is a spectrum. Model developers release different components, use different licenses, and deploy to users via download, application, or online service. 'Open' or 'closed' poorly predicts risk because the use and user matter more. [1]

WHY OPEN MODELS MATTER

1. **Competition and entry.** Open weights let entrants skip most of the cost of training a foundation model, so firms compete on customization, applications, safety, privacy, and services. This can discipline large model providers more than regulation. [5]
2. **Choice and independence.** Users adapt models, keep sensitive data on in-house computers, and can switch vendors. If only a few closed models serve everyone, only a few companies decide which tools exist and how they can be used.
3. **Safety and resilience.** Open weights let outside researchers test and verify developer claims and repair and improve models.
4. **American leadership.** Retreat from open release cedes ground to capable foreign models, and to their biases and limits. [7]

WHO SHOULD PAY WHEN PEOPLE ABUSE OPENNESS' FREEDOM?

Anyone who downloads an open model can modify it and pass it on. These benefits also mean that built-in safeguards may not survive such changes. Some will abuse that freedom. Closed models get misused too, and a free society should not cripple a general-purpose tool for everyone to stop a few. The developer of an open weight model can't control the party that picks the application, supplies the data, or distributes illegal content. Developers should answer for their own unlawful conduct, but not for what unaffiliated third parties do with a general-purpose model. [4][5]

TO KEEP OPEN WEIGHT MODELS VIABLE, LEGISLATION SHOULD

- **Identify a concrete harm.** Define the injury, the people affected, the conduct, and the evidence required. Fraud, unlawful discrimination, nonconsensual intimate imagery, physical injury, and property damage make clearer targets than vague references to AI risk.
- **Apply existing law first.** Write a new AI-specific rule only after finding a gap in existing fraud, civil-rights, tort, consumer-protection, criminal, or privacy law. Often increased enforcement of an existing law would better fix the issue than would creating new law. If a bill's sponsor cannot identify a gap in existing law, that's a solution searching for a problem. [2]
- **Put duties where control sits, and keep them feasible.** Obligations should fall on those with knowledge and control. No developer can monitor or control open weights once released. Write light-touch safe harbors, drafted so a voluntary standard never hardens into a compliance floor only the largest labs can meet. [4]
- **Preserve a national market.** Preempt conflicting state rules for model development, and leave states their authority over local harmful uses. A patchwork burdens startups and open projects most. [6]
- **Build the American open-model ecosystem.** Direct the Department of Commerce to identify barriers, expand evaluation resources, back voluntary benchmarks, and pilot federal use of qualified American open models. Open federal datasets and funded research to developers. [7]
- **Distinguish the model from the service.** Downloaded weights running on American hardware raise fewer privacy and security concerns than an app that forwards every prompt to a foreign company. Where national origin warrants scrutiny, government labs could evaluate and publicize model security, capability, accuracy, and cost. [3][7]
- **Protect speech and lawful computing.** Courts have treated code as protected speech, so restricting weight release raises serious First Amendment problems, and mandating government-preferred answers raises sharper ones. Restrictions on computing should be narrowly tailored. [1][6]

BILLS WITH THESE RED FLAGS THREATEN OPEN WEIGHT MODELS

- Licensing, preapproval, or FDA-style review, and compute or parameter thresholds standing in for evidence of risk.
- “Capable of,” “could be used to,” “reasonably foreseeable misuse.” These target the tool rather than the conduct.
- Publication as the trigger: “makes available,” “publishes,” “releases.” Release is not an injury.
- Duties to monitor, control, or “ensure” downstream use, which no decentralized project can perform.
- Private rights of action untethered to a demonstrated injury, and compelled ideological conditioning.

BILL-REVIEW CHECKLIST

1. Does the bill name a specific, demonstrated harm and say why existing law falls short? **If not, ask for both before markup.** [2]
2. Does each duty fall on someone with the knowledge and control to perform it? **If it lands on the developer, probably move it downstream.**
3. Could a university lab or a two-person startup comply? **If only the largest labs can, the bill entrenches them.**
4. Does the bill reach harmful conduct, or the release of a tool? **If release triggers the duty, retarget it at use.**
5. Is the concern the model or the service? **If data flowing abroad is the worry, regulate the service.**
6. Does the analysis weigh foregone benefits, lawful speech and research, and the risk of a state patchwork? **If not, ask for that accounting.**

SELECTED ABUNDANCE INSTITUTE SOURCES

[1] [NTIA open-weights comment \(2024\)](#)

[4] [AI Action Plan comment \(2025\)](#)

[7] [House technology-leadership testimony \(2026\)](#)

[2] [Evaluating AI Policy Proposals \(2025\)](#)

[5] [House competition testimony \(2025\)](#)

[3] [DeepSeek policymaker's guide \(2025\)](#)

[6] [Federal tools to address state overreach \(2025\)](#)