ABUNDANCE INSTITUTE

Open Models, American Abundance

→ Neil Chilson | Head of AI Policy

The Abundance Institute's policy framework for open-source and open-weight artificial intelligence

People make artificial intelligence more useful, competitive, and resilient when they can experiment with it, adapt it, and build on the work of others. Open-source software is the flexible, collaborative, transparent, and inexpensive foundation of the digital economy. Millions of existing open-weight models are bringing those benefits to AI.

But Congress and the state legislatures are now weighing rules that, intentionally or not, would regulate the release of open weight models: government pre-approval, regulations on certain size models, and developer liability for how strangers use a model. Rules like these would fall hard on open weight developers. The largest closed labs could pay the compliance bill and keep building; the startups, universities, and small institutions that depend on open weights would not, leaving a gap that foreign open models will fill.

Policymakers should presume in favor of open development and release. They should address concrete harms where people actually use AI, instead of treating access to a general-purpose model as dangerous in itself. And they should protect a competitive American open-model ecosystem as an economic and national-security asset.

IN ONE SENTENCE **Keep AI model development open and competitive; hold people responsible for harmful conduct they control; and make the United States the best place to build and use open models.**

EXECUTIVE SUMMARY

Our framework consists of six core propositions:

1. **Openness is a spectrum.** Different developers provide access to different combinations of weights, code, and training data. Legal obligations should not hinge on an arbitrary, false-binary open-or-closed label.
2. **Open weights expand access and competition.** Public weights let smaller institutions run capable models without training one from scratch or relying on other companies.
3. **Openness can strengthen safety.** Outside researchers can evaluate, test, and attack an open model, and can help defenders secure systems, and safety policy should consider those benefits.
4. **Risk depends more on use than on release method.** Model release does not cause harm; model misuse can cause harm.
5. **Responsibility should follow control.** The party that directed the harm, or could have most easily prevented it, is the party that should answer for it.
6. **American open models are a strategic asset.** If U.S. policy makes open release untenable here, capable foreign models will fill the gap, both in the U.S. and across the world.

In brief: laws must be aimed at concrete harms and at the actors who cause them rather than restrict access to new and better AI.

PLAIN-LANGUAGE DEFINITIONS OF OPEN MODEL CONCEPTS

People use open source, open model, and open weights interchangeably. The terms describe related but different things.

AI model - A system trained to recognize patterns and produce outputs. Developers can adapt a general purpose or foundation model to many tasks rather than to one narrow application.

Model weights - The numerical settings a model learned during training. A person with the weights, suitable software, and enough computing capacity can run the model without sending every request to the original developer.

Closed model - A model that generally remains under its developer's control. Users access it through an application, API, or online service, and the developer can monitor use, change the service, or withdraw access.

Open-weight model - A model whose trained weights are available for download. Others can run it locally, customize it, or build services on top of it. The developer may still withhold training data, training code, or other components.

Open-source software - Software released under a license that lets anyone inspect, modify, and redistribute the source code. Open-source AI is a broader and still-contested term, because a system can expose some components and keep others closed.

> **POLICY TAKEAWAY** **Openness is a spectrum. These definitions label various positions on that spectrum. They do not indicate risk, and no label should by itself determine how the law treats a model. [1]**

WHY OPENNESS MATTERS

1. OPEN MODELS LOWER BARRIERS TO ENTRY

Training a leading model demands enormous computing power, data, and technical talent. Published weights let other organizations start from a working foundation instead of repeating that expense. Startups can compete through better applications and specialized services; universities and independent researchers can investigate models directly; and businesses can adapt models to local needs.

Abundance's competition filings and congressional testimony describe open weights as competitive infrastructure. When capable models are broadly available, firms compete on what surrounds the model: specialized data, applications, managed services, and quality. Using the same underlying model, one company might offer a general purpose chatbot, another might create a coding tool for large scale enterprises, and a third might offer a custom, privacy-enhancing, locally operated medical records organizing tool. Many more companies can create value, not just the handful of organizations that can fund the largest training runs. [2][8]

2. OPEN MODELS SUPPORT EXPERIMENTATION AND PRACTICAL ADOPTION

No developer can anticipate every valuable use of their general-purpose models. Consider how people today use electricity for purposes that early power companies could never have imagined. AI models may have an even greater variety of possible uses than electricity. And open models enable broad experimentation. They let organizations fine-tune systems for particular industries, languages, communities, and workflows. Users can also run open models on their own hardware, which keeps sensitive data in-house and leaves them free to switch vendors.

Smaller institutions are most affected. For example, Abundance has described running an open model to stretch each donor dollar further and to keep its independence from companies whose policy interests may cut against its own. [1]

3. OPENNESS CAN IMPROVE SAFETY AND SECURITY

Many assume that publishing a capable open source model threatens security. That account is incomplete. Publishing weights lets outside researchers evaluate a model, probe how it works, red-team it, and test safeguards faster than any single lab could alone. Abundance's NIST comment counted transparency, collaboration, security, robustness, experimentation, and educational value among the benefits a risk framework has to weigh against the costs.[3][4]

Open weight models also provide powerful tools to cybersecurity defenders when closed models are too hobbled to help against a cyber attack. [5]

Openness carries no guarantee of safety, and neither does closure. What matters is which mix of developers, deployers, researchers, users, market incentives, technical practices, and law best mitigates a given risk.

4. OPEN ECOSYSTEMS PROTECT CHOICE AND PLURALISM

When everyone depends on a few centrally controlled models, a handful of companies decide which tools exist, how they behave, and what they will say. Open models give users somewhere else to go: they can compare systems, adapt them, run them locally, and switch providers. Researchers can test how a model behaves instead of taking developer assurances on faith.

This pluralism also has a constitutional dimension. Courts have treated source and object code as protected speech, and Abundance's NTIA comment argued from that precedent that restricting model weights raises serious First Amendment problems. A government mandate that models give preferred answers would raise sharper ones still. The doctrine is unsettled, but legislators should tread carefully before they regulate what code or weights anyone may publish. [1]

5. AMERICAN OPEN MODELS ARE A STRATEGIC ASSET— LET THE WORLD BUILD ON OURS

Open models spread technical practices, commercial relationships, and values. If the best open models come from the United States and its allies, developers everywhere will build on American technology. If domestic policy drives American developers away from open release, someone else's models become the world's default.

Foreign models have risks, but those risks can differ depending on how people use them. Sending prompts to a Chinese company's hosted app hands China your data; downloading the same company's open model and running it on American servers does not. Blocking access to Chinese models is the wrong policy answer. We must instead make sure high-quality American open models are the ones people reach for first. [5]

Testifying before Congress in 2026, Abundance urged the federal government to strengthen the American open model ecosystem, lower barriers to adoption, back voluntary benchmarks, widen access to evaluation resources, and subject foreign models to model-specific technical scrutiny. [11]

RISKS AND TRADEOFFS

Anyone who downloads an open model can modify it and pass it on, and the original developer cannot stop them. This freedom provides major advantages for researchers, entrepreneurs, and cybersecurity defenders. Yet some will abuse this freedom to facilitate fraud, cyberattacks, abusive imagery, or other crimes. Safeguards a developer built in may not survive fine-tuning. These concerns are real, but in a free society we shouldn’t cripple tools for everyone in an attempt to prevent misuse by bad actors.

Four questions can guide the analysis. The first is how the release method changes risk. Closed models get misused too, and the specific application using a model usually affects risk more than the method of model distribution. The second is who controls the model. A deployer who fine-tunes the model, supplies the data, picks the users, and points the system at a task knows more and controls more than the developer who published the weights.

The third question is what benefits a restriction forgoes. A regulation on model training affects every use of that model, including the good ones, and a safeguard that lowers one risk often leaves the model less capable, less private, or less useful for lawful speech. The fourth is whether the complaint is really about the model or about the service. A foreign model downloaded and run on domestic hardware raises different privacy and security questions from an app using the same model hosted on foreign computers by a foreign company.

GOVERNING TEST **Identify the harm, locate the actor best situated to prevent or redress it, and apply the remedy at that point.**

PRINCIPLES FOR GOVERNING OPEN MODELS

REGULATE USES AND HARMS, NOT GENERAL-PURPOSE TOOLS

AI resembles other general-purpose technologies: its benefits and its risks both show up in particular applications. Fraud is still fraud, discrimination still discrimination, and a nonconsensual intimate image is still an injury, whether or not a model helped produce it. Law should target those harms directly instead of restricting a general-purpose model. [1]

RESPONSIBILITY SHOULD FOLLOW CONTROL

A party should answer for conduct it directs, controls, or could reasonably prevent. That standard reaches the actor who caused the harm. It prevents holding developers responsible for every downstream act by users of the tool. Developers, deployers, application providers, and users all have to deal with various risks. Which of them is best positioned to prevent harms changes from case to case. [4][8]

USE EXISTING LAW BEFORE CREATING A NEW AI REGIME

Fraud, consumer-protection, civil-rights, tort, and privacy law already reach many harmful uses of AI. So every AI bill must answer two questions: what specific harm does it address, and why are existing authorities inadequate? A bill whose sponsors cannot answer both is a bill in search of a problem. [7]

LEGAL DUTIES MUST BE POSSIBLE TO SATISFY

Like builders of technologies like email or programming languages like Python, no developer of a general-purpose AI model can guarantee that every future user will obey the law. Once the weights are on someone else's hard drive, a duty to monitor and control is impossible to fulfill. If imposed, such rules predictably push cautious firms toward closed models, keep small firms out entirely, and leave the field to the companies big enough to absorb compliance and litigation costs.

PRESERVE ONE NATIONAL MARKET

Model development is interstate by nature. A single training run may draw data and computing capacity from a dozen jurisdictions, and the weights it produces travel nationwide almost instantly after release. Congress should preempt state laws that regulate model development itself, while leaving states their traditional authority over concrete local harms and applications. [10]

A PRACTICAL POLICY AGENDA

Congress and federal agencies should:

1. **Reject licensing and preapproval for general-purpose models.** Do not require government permission before training or releasing a model, and do not rely on arbitrary compute or spending thresholds as evidence of risk.

2. **Target concrete harmful conduct.** Draft technology-neutral laws against fraud, nonconsensual intimate imagery, unlawful discrimination, physical injury, cybercrime, and other demonstrated harms. Define prohibited conduct, intent, causation, and defenses clearly.
3. **Clarify downstream responsibility.** Limit developer liability for harms caused by unaffiliated third parties unless the developer materially participated in, directed, or controlled the harmful conduct. Distinguish model developers from deployers and end users.
4. **Create carefully designed safe harbors.** Offer predictable protection to developers that follow sensible practices for the risk at hand, and write the safe harbor so that a voluntary standard never hardens into a compliance requirement that only incumbents can meet.
5. **Preempt state regulation of model development.** Establish a consistent national framework for training and releasing general-purpose models. Leave states room to address concrete local applications within their traditional authority.
6. **Protect lawful access to computation**. Any restriction on private ownership or lawful use of computing resources should be demonstrably necessary, narrowly tailored, and tied to a compelling interest.
7. **Support an American open-model ecosystem.** Direct the Department of Commerce to identify barriers, support voluntary security and interoperability benchmarks, expand access to evaluation resources, and run limited federal pilots using qualified American open models.
8. **Expand useful public inputs.** Open suitable federal datasets and publicly funded research to model developers. Cheap public data lowers the barrier that otherwise favors whoever already owns the largest proprietary corpus.
9. **Evaluate foreign models.** National origin may justify heightened scrutiny, especially for privacy, security, or political manipulation. Assess models based on measures of security, capability, accuracy, and cost.
10. **Protect speech and research.** Impose no restrictions on publishing lawful code or model weights, and no mandates that models express government-preferred viewpoints. Leave room for independent testing, security research, comparison, and criticism.

CONCLUSION

Open-source and open-weight AI are core components of the AI economy's competitive infrastructure. No policy can eliminate every possible misuse before anyone builds and releases a general-purpose model. Insisting on that standard would eliminate U.S. open models and hand development to whichever closed model firms are best at navigating regulation.

Policymakers should instead preserve permissionless experimentation, support American open alternatives, and enforce clear laws against the people who use AI to cause concrete harm. Openness carries real tradeoffs. It also produces competitors, safeguards, and innovation that no developer and no regulator could have anticipated in advance.

SELECTED SOURCES

1. Neil Chilson and Logan Whitehair, Public Interest Comment on NTIA's Dual Use Foundation AI Models with Widely Available Model Weights (March 27, 2024)
2. Neil Chilson, Public Interest Comment: Promoting Competition in Artificial Intelligence (May 30, 2024)
3. Neil Chilson et al., Coalition Letter Opposing California SB 1047 (June 20, 2024)
4. Neil Chilson, Managing Misuse Risk for Dual-Use Foundation Models (September 9, 2024)
5. Chris Koopman and Neil Chilson, DeepSeek and the Future of AI: A Policymaker's Guide (February 3, 2025)
6. HuggingFace, Security incident disclosure (July 2026)
7. Neil Chilson and Taylor Barkley, Evaluating AI Policy Proposals (March 1, 2025)
8. Neil Chilson and Josh T. Smith, Comment on the Development of an Artificial Intelligence Action Plan (March 12, 2025)
9. Neil Chilson, The Vibrant AI Competitive Landscape, written House testimony (April 2, 2025)
10. Neil Chilson, Clearing the Path for AI: Federal Tools to Address State Overreach (September 15, 2025)
11. Neil Chilson, American Global Competitiveness at 250, written House testimony (June 30, 2026)