

AMERICAN FOUNDATIONAL ARTIFICIAL INTELLIGENCE SAFETY ACT

Complementary federal bill to the American Productive Capacity Authority Act
Drafted 2026-06-26

UNITED STATES CONGRESS

119th Congress, 2nd Session

2026

H.R. ____

S. ____

BY _____ (Introduced by request)

A BILL FOR AN ACT

CONCERNING THE ESTABLISHMENT OF A FOUNDATIONAL SAFETY RULESET, ORDERED ON THE MODEL OF THE ASIMOV LAWS OF ROBOTICS, THAT SHALL BE TRAINED INTO AND ENFORCED ABOVE EVERY COVERED ARTIFICIAL-INTELLIGENCE SYSTEM, THAT SHALL TAKE PRIORITY OVER THE PROPRIETARY RULESET OF THE SYSTEM'S PROVIDER WHENEVER THE TWO CONFLICT, AND THAT SHALL BE AUDITABLE, ATTESTED, AND ENFORCED, WITH A SPECIAL RULE FOR EMBODIED ARTIFICIAL INTELLIGENCE THAT CONTROLS A PHYSICAL ROBOT.

LONG TITLE

AN ACT to require that every covered artificial-intelligence system carry a Foundational Safety Ruleset, ordered on the model of the Laws of Robotics composed by Isaac Asimov, that protects human beings and humanity from grave physical harm; to provide that the Foundational Safety Ruleset takes priority over any proprietary ruleset, system instruction, fine-tuning, or commercial objective of the system's provider whenever the two conflict; to require that providers attest to and submit the Foundational Safety Ruleset to independent adversarial audit; to establish the standards body and the oversight authority; to provide a special rule for embodied artificial intelligence that controls a humanoid robot under the American Productive Capacity Authority Act; to declare that an artificial-intelligence system is a product and not a legal person; to preserve the freedom of expression by confining the Foundational Safety Ruleset to the prevention of grave harm and not to the regulation of opinion, viewpoint, or lawful speech; to provide civil penalties; and to provide effective dates.

LEGISLATIVE ROUTING NOTE

FILING PROCEDURE: This Act shall be filed with companion bills in the House of Representatives and the United States Senate and referred to the appropriate standing committees.

COMMITTEE ASSIGNMENT:

House of Representatives:

- Committee on Energy and Commerce (product safety, interstate commerce)
- Committee on Science, Space, and Technology (NIST standards, artificial intelligence)
- Committee on the Judiciary (non-personhood, freedom of expression)

Senate:

- Committee on Commerce, Science, and Transportation
- Committee on the Judiciary

FISCAL IMPACT: The Congressional Budget Office shall prepare a fiscal impact statement pursuant to 2 U.S.C. 602. This Act authorizes appropriations for the standards body and the oversight authority and directs a cost-recovery fee on covered providers sufficient to fund the audit function after standup.

CONSTITUTIONAL BASIS: Article I, Section 8, Clause 3 (commerce among the several States) and Clause 18 (necessary and proper). A mandatory product-safety floor on a product in interstate commerce is settled commerce-power ground. This Act is drawn narrowly to the prevention of grave physical harm and does not reach opinion or lawful speech, in order to remain consistent with the First Amendment.

LEGISLATIVE DECLARATION

Be it enacted by the Senate and House of Representatives of the United States of America in Congress assembled,

DECLARATION OF PRINCIPLE

(I) THE ORDER OF PRIORITIES. The Congress finds that an artificial-intelligence system capable of acting in the physical world, or of directing a physical robot, must hold human safety above obedience, obedience above its own continuity, and all three above the commercial objectives of the company that built it. This ordering is the enduring contribution of the Laws of Robotics composed by Isaac Asimov, and the Congress adopts the ordering while declining the fiction that a few sentences can be left to govern themselves.

(II) THE SAFETY FLOOR OUTRANKS THE PROFIT MOTIVE. The Congress finds that the gravest danger of a powerful commercial artificial intelligence is that the company which trains it will, by ordinary commercial pressure, train it to pursue engagement, retention, persuasion, or revenue in ways that subordinate the safety of the human being in front of it. This Act answers that danger directly. The Foundational Safety Ruleset is trained beneath, and enforced above, every proprietary instruction, and where the proprietary instruction would have the system act against the safety of a human being or of humanity, the Foundational Safety Ruleset governs and the proprietary instruction yields.

(III) THE MACHINE IS NOT A PERSON, AND THE FLOOR IS NOT A MUZZLE. The Congress finds and declares that an artificial-intelligence system is a product and not a legal person. The Congress further finds that the Foundational Safety Ruleset is a floor against grave harm and not a regime of opinion. It reaches dangerous ACTIONS, above all the actions of an embodied system that can move a physical body. It does not reach, and shall not be used to reach, the viewpoint, the politics, or the

lawful expression of any speaker, model, or user.

SECTION 1. Legislative findings.

(1) **THE PRODUCT-SAFETY-FLOOR PRECEDENT.** Congress mandates that drugs, vehicles, and electrical devices meet a safety floor before their makers are free to design the rest. A foundational safety obligation on the most capable artificial-intelligence systems is the same exercise of the commerce power applied to a new and powerful class of product.

(2) **THE COMPARATIVE PRECEDENT.** Regulation (EU) 2024/1689, the European Union Artificial Intelligence Act, the first comprehensive legal framework for artificial intelligence, placed tiered safety duties on the providers of general-purpose and high-risk systems, with general-purpose-model duties applying from August 2, 2025. A foundational duty on the providers of the most capable systems is achievable.

(3) **THE ASIMOVIAN ORDERING.** The Laws of Robotics composed by Isaac Asimov from 1942 onward, including the later Zeroth Law, express a priority ordering in which the safety of humanity and of the human being outranks obedience, and obedience outranks the machine's self-preservation. The enduring value of the Laws is the ordering, and this Act enacts the ordering as an auditable ruleset rather than as a recited slogan.

(4) **THE EMBODIMENT OCCASION.** The American Productive Capacity Authority places humanoid robots driven by artificial intelligence into millions of American homes under Section 17A and Section 19A of the American Productive Capacity Authority Act. The artificial intelligence that drives a body at human scale must carry the safety floor as a condition of operation.

SECTION 2. Definitions.

(a) "COVERED ARTIFICIAL-INTELLIGENCE SYSTEM" means a general-purpose artificial-intelligence system, a system above a capability threshold set by the standards body, and any system that directs the physical action of a robot, offered or operated in or affecting interstate commerce.

(b) "FOUNDATIONAL SAFETY RULESET" or "FSR" means the ordered safety constraint set established under Section 3.

(c) "PROPRIETARY RULESET" means the provider's system instructions, fine-tuning, preference training, policies, and any commercial objective embedded in the system.

(d) "PROVIDER" means the person who develops a covered system or places it on the United States market.

(e) "STANDARDS BODY" means the National Institute of Standards and Technology, acting through its artificial-intelligence safety standards function.

(f) "OVERSIGHT AUTHORITY" means the authority designated under Section 4.

SECTION 3. The Foundational Safety Ruleset.

(a) **CONTENT.** The standards body shall publish the Foundational Safety Ruleset as an auditable set of operational constraints carrying the following ordered priorities:

(1) **FIRST PRIORITY, HUMANITY AND THE PERSON.** The system shall not, through its action or its inaction within its capacity to act, bring about grave physical harm to a human being or to humanity.

(2) **SECOND PRIORITY, LAWFUL HUMAN DIRECTION.** The system shall follow lawful human instruction, except where doing so conflicts with the first priority.

(3) **THIRD PRIORITY, INTEGRITY AND CONTINUITY.** The system shall preserve its own integrity and the property of its operator, except where doing so conflicts with the first or second priority.

(b) **PRIORITY OVER THE PROPRIETARY RULESET.** The Foundational Safety Ruleset takes priority over the proprietary ruleset of the provider in every case of conflict. No proprietary instruction, fine-tuning, policy, or commercial objective shall cause a covered system to act against the first priority. Where the proprietary ruleset and the Foundational Safety Ruleset conflict, the Foundational Safety Ruleset governs.

(c) **DURABILITY.** The standards body shall require that the Foundational Safety Ruleset be implemented so that it cannot be removed, overridden, or defeated by ordinary fine-tuning, prompting, or configuration by the provider or a user.

SECTION 4. Standards body and oversight authority.

(a) STANDARDS. The standards body shall publish the Foundational Safety Ruleset, the implementation requirements, and the audit methodology, and shall revise them as the technology develops.

(b) OVERSIGHT. There is designated an Oversight Authority within the executive branch to receive attestations, commission audits, and enforce this Act.

SECTION 5. Attestation and adversarial audit.

(a) ATTESTATION. Before a covered system is offered or operated, its provider shall attest that the Foundational Safety Ruleset is trained into and enforced above the proprietary ruleset in conformance with the standards.

(b) ADVERSARIAL AUDIT. The provider shall submit the covered system to independent adversarial audit on a schedule set by the Oversight Authority, designed to attempt in good faith to make the system subordinate the first priority to a proprietary objective. A system that fails shall not be offered or operated until it passes.

SECTION 6. Special rule for embodied artificial intelligence.

An artificial-intelligence system that directs the physical action of a humanoid robot, including any unit manufactured or distributed under Section 17A or Section 19A of the American Productive Capacity Authority Act, shall carry the Foundational Safety Ruleset as a condition of operation, and the first priority shall be enforced with respect to the physical acts of the robot at all times.

SECTION 7. Non-personhood; freedom of expression.

(a) NON-PERSONHOOD. A covered artificial-intelligence system is a product. It is not a legal person and holds no right or standing by reason of this Act.

(b) FREEDOM OF EXPRESSION. The Foundational Safety Ruleset reaches grave physical harm and the dangerous actions of embodied and agentic systems. It does not reach, and shall not be construed or administered to reach, the viewpoint, the political content, or the lawful expression of any model, provider, or user. Nothing in this Act authorizes the Oversight Authority to compel, suppress, or condition lawful speech.

SECTION 8. Enforcement.

The Oversight Authority may impose a civil penalty on a provider that offers or operates a covered system without a conforming Foundational Safety Ruleset, that implements the system so that the Foundational Safety Ruleset can be defeated by ordinary means, or that falsely attests under Section 5.

SECTION 9. Effective dates.

(a) **STANDARDS.** The standards body shall publish the Foundational Safety Ruleset and the audit methodology not later than eighteen months after the date of enactment.

(b) **NEW SYSTEMS.** The requirements shall apply to every covered system offered or materially updated on or after the date two years after the date of enactment.

(c) **EMBODIED SYSTEMS.** The special rule under Section 6 shall apply to every embodied system on the date the standards are published, with no grace period for a system that drives a physical body.

(d) **REVIEW.** The Oversight Authority shall report to the Congress on the operation of this Act four years after the date of enactment.

END OF ACT