

(7) Coordination.—If the Commission, the Antitrust Division, or the United States has instituted a civil action for the same alleged violation, the court may stay, consolidate, coordinate, or otherwise manage an action brought under paragraph (4) to avoid unnecessary duplication, but an action under paragraph (4) shall not be barred solely because of the pendency of a civil action brought by the Commission, the Antitrust Division, or the United States.

(8) Rules.—Not later than January 1, 2027, the Commission, in consultation with the Antitrust Division, shall issue interim final rules to carry out this section and shall thereafter promulgate final rules after notice and comment.

(j) Rule of construction.—

(1) No duty to disclose model weights.—Nothing in this section shall be construed to require a foundation model provider to disclose, publish, transfer, license, or make available model weights, source code, training data, system prompts, cybersecurity information, safety mitigations, trade secrets, or other confidential business information, except to the Commission or the Antitrust Division as necessary to enforce this section.

(2) No duty to provide unsafe or unlawful access.—Nothing in this section shall be construed to require a foundation model provider to provide, continue, or restore access to a model access service where such access would violate Federal, State, or local law, sanctions, export controls, a court order, or another lawful order issued pursuant to statutory authority, or would create a material and reasonably foreseeable risk of severe harm that cannot be mitigated through a narrower measure.

(3) No limitation on safety rules.—Nothing in this section shall be construed to prohibit a foundation model provider from adopting and enforcing bona fide, generally applicable, publicly disclosed, and consistently applied safety, security, abuse prevention, or legal compliance policies.

(4) No limitation on antitrust laws.—Nothing in this section shall be construed to limit the authority of the Commission, the Antitrust Division, an attorney general of a State, or any other person under the antitrust laws.

(5) No limitation on other law.—Nothing in this section shall be construed to limit any stricter duty, liability, remedy, penalty, or prohibition otherwise available under Federal or State law.

Subtitle F—Regulating Artificial Intelligence

SEC. 581. DEFINITIONS.

(a) Definitions.—In this subtitle:

(1) Advertising system.—The term "advertising system" means any service or system that facilitates the dissemination of advertising to users, viewers, listeners, or other members of the public, and includes:

(A) a digital advertising platform, exchange, network, marketplace, or intermediary;

(B) a broadcast, cable, satellite, streaming, audio, audiovisual, or connected-device advertising service;

(C) an advertising sales, brokerage, placement, targeting, measurement, or optimization service; and

(D) any other service or system that sells, places, targets, delivers, measures, optimizes, brokers, or otherwise facilitates advertising.

(2) Artificial intelligence; AI.—The terms "artificial intelligence" and "AI" mean a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, decisions, or generate content influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to—

(A) perceive real and virtual environments;

(B) abstract such perceptions into models through analysis in an automated manner; and

(C) use model inference to formulate options or content for information or action.

(3) Artificial intelligence chatbot; chatbot.—The terms "artificial intelligence chatbot" and "chatbot" mean a generative artificial intelligence system that—

(A) produces new expressive content not fully predetermined by the developer;

(B) accepts open-ended natural-language or multimodal user input; and

(C) produces adaptive or context-responsive output.

(4) Clear and conspicuous.—The term "clear and conspicuous", with respect to a disclosure, means that the disclosure meets the following criteria:

(A) for any content that is solely visual or solely audible, the disclosure shall be made through the same means through which the content is presented;

(B) for any content that is both visual and audible, the disclosure shall be visual and audible;

(C) a visual disclosure, by its size, contrast, location, the length of time it appears, and other characteristics, shall stand out from any accompanying text or other visual elements so that the disclosure is easily noticed, read, and understood;

(D) an audible disclosure shall be delivered in a volume, speed, and cadence sufficient for a reasonable person to easily hear and understand the disclosure;

(E) the disclosure shall not be avoidable;

(F) the disclosure shall use diction and syntax understandable to a reasonable person;

(G) the disclosure shall not be contradicted, mitigated, or rendered inconsistent by any other part of the communication; and

(H) the disclosure shall satisfy any additional criteria prescribed by the Commission by rule.

(5) Commission.—The term "Commission" means the Federal Trade Commission.

(6) Digital person.—The term "digital person" means a public-facing artificial intelligence system that is presented to users as a persistent or recurring simulated persona, and includes a chatbot when so presented.

(7) Digital replica.—The term "digital replica" means a computer-generated, algorithmically generated, or artificial intelligence-generated voice, image, likeness, avatar, style, performance, conversational manner, or other identifiable simulation of a natural person.

(8) First interaction.—The term "first interaction" means, for a calendar day, the first direct exchange, communication, contact, or other point at which a user encounters, engages with, or is exposed to an artificial intelligence system, digital person, or synthetic media in a context regulated under this subtitle, including any rules issued pursuant thereto.

(9) High-risk context.—The term "high-risk context" means any context involving financial products or services, insurance, health care, housing, employment, education, legal services, debt collection, benefits navigation, government-related services, identity verification, age verification, elections, fundraising, or any other context designated by the Commission by rule as presenting a heightened risk of fraud, deception, exploitation, or material consumer harm.

(10) National Institute of Standards and Technology; NIST.—The terms "National Institute of Standards and Technology" and "NIST" mean the National Institute of Standards and Technology of the Department of Commerce.

(11) Persistent disclosure.—The term "persistent disclosure" means a disclosure that, beginning with a first interaction, remains continuously visible, intermittently audible, or otherwise reasonably perceptible throughout the period in which a user is interacting with, viewing, listening to, or otherwise engaging with a digital person, synthetic media, or other artificial intelligence output for which persistent disclosure is required under this subtitle.

(12) Person.—The term "person" has the meaning given such term in section 1 of the Clayton Act ([15 U.S.C. 12](#)).

(13) Provenance data.—The term "provenance data" means information attached to, embedded in, associated with, or otherwise linked to content that identifies or describes whether the content was generated or materially altered by artificial intelligence, the source or provider of the content, the time or manner of generation or alteration, and any other information required by this subtitle.

(14) Public-facing artificial intelligence system.—The term "public-facing artificial intelligence system" means an artificial intelligence system that is made available to the public for communication, interaction, content creation, performance, recommendation, or other expressive or commercial activity.

(15) Secretary.—The term "Secretary" means the Secretary of Commerce.

(16) Synthetic media.—The term "synthetic media" means any image, audio recording, audiovisual work, text, likeness, performance, or other expressive content that has been generated, materially altered, or materially manipulated in whole or in part by artificial intelligence. Such term includes digital replicas.

(17) Synthetic performer.—

(A) In general.—The term "synthetic performer" means—

(i) a visual, auditory, audiovisual, or interactive representation;

(ii) generated, materially altered, or materially simulated by computational, algorithmic, or artificial intelligence means; and

(iii) likely to cause a reasonable person to believe that such representation is a portrayal or performance of a natural person.

(B) Exclusion.—The term "synthetic performer" does not include editing, lighting correction, makeup, color grading, dubbing, captioning, or similar visual or audio cleanup of a portrayal or performance of a natural person, if such activity does not materially change the identity, speech, conduct, appearance, or performance of such person.

(18) Tribal government.—The term "Tribal government" has the meaning given the term "Indian Tribe" in section 4 of the Indian Self-Determination and Education Assistance Act ([25 U.S.C. 5304](#)).

SEC. 582. FRONTIER-MODEL SECURITY SAFEGUARDS.

(a) Definitions.—In this section:

(1) Adverse AI incident.—The term "adverse AI incident" means any known or reasonably suspected event in which an artificial intelligence model—

(A) materially contributes to, materially enables, or materially increases the risk of death, serious bodily injury, serious property damage, substantial critical infrastructure disruption, substantial cybersecurity compromise, or other severe public harm;

(B) materially enables the design, acquisition, or use of a biological, chemical, radiological, or nuclear weapon;

(C) materially enables autonomous cyber operations capable of causing substantial disruption to critical infrastructure, government systems, or major economic activity; or

(D) materially defeats or evades implemented safeguards.

(2) Artificial intelligence blue-teaming.—The term "artificial intelligence blue-teaming" means an effort to conduct operational vulnerability evaluations and provide mitigation techniques to entities who have a need for an independent technical review of the security posture of an artificial intelligence system.

(3) Artificial intelligence red-teaming.—The term "artificial intelligence red-teaming" means structured adversarial testing efforts of an artificial intelligence system to identify risks, flaws,

and vulnerabilities of the artificial intelligence system, such as harmful outputs from the system, unforeseen or undesirable system behaviors, limitations, or potential risks associated with the misuse of the system.

(4) Cloud compute provider.—The term "cloud compute provider" means any person that offers, provides, or makes available, by means of remote access, cloud or other network-accessible computing services, including processing, storage, networking, model-hosting, or related computational resources.

(5) Deploy; deployment.—The terms "deploy" and "deployment" mean to use, or the use of, a frontier model in commerce, by contract, by license, by application programming interface, by hosted service, by incorporation into another system, or by any other means by which the model is made available to another person.

(6) Developer.—The term "developer", with respect to a frontier model, means any person that—

(A) develops, trains, or materially fine-tunes the model; or

(B) controls or is responsible for the release, deployment, or continued availability of the model.

(7) Foreign person.—The term "foreign person" means any individual who is not a citizen or national of the United States and is not an individual lawfully admitted for permanent residence in the United States, and any corporation, partnership, association, trust, estate, or other entity that is not organized under the laws of the United States or of any jurisdiction within the United States.

(8) Foreign reseller.—The term "foreign reseller" means any person outside the United States that resells, brokers, leases, sublicenses, or otherwise makes available cloud compute services of a cloud compute provider to another person.

(9) Frontier model.—The term "frontier model" means an artificial intelligence model that—

(A) was trained using more than 10^{26} integer or floating-point operations; or

(B) is designated by the Secretary under subsection (b)(2).

(10) Frontier-model-relevant compute.—The term "frontier-model-relevant compute" means compute, storage, networking, or related computational resources reasonably capable of training a frontier model or otherwise meeting a threshold prescribed by the Secretary by rule under this section.

(11) Open-weight release.—The term "open-weight release" means the distribution, publication, transfer, licensing, or other making available of the weights or functionally equivalent model parameters of a frontier model in a manner that permits another person to run, fine-tune, modify, or redistribute the model without the continuing operational control of the original developer.

(12) Operate.—The term "operate" means to cause an artificial intelligence model to perform computation or generate outputs, whether or not in commerce.

(b) Frontier model evaluation and security program.—

(1) In general.—The Secretary, acting through the National Institute of Standards and Technology, shall establish and administer a frontier model evaluation and security program for frontier models.

(2) Designation authority.—The Secretary may designate an artificial intelligence model as a frontier model if, after notice to the developer and an opportunity to respond, the Secretary determines that the model exhibits, or could readily be adapted to exhibit, capabilities that pose a substantial risk of materially enabling an adverse AI incident.

(3) Duties.—The program established under paragraph (1) shall—

(A) establish evaluation protocols, artificial intelligence red-teaming requirements, reporting requirements, security requirements, and recordkeeping requirements for frontier models;

(B) establish reliable channels and procedures for receiving, documenting, triaging, escalating, and acting upon reports of adverse AI incidents and material security incidents involving frontier models, whether submitted by developers, deployers, researchers, affected persons, artificial intelligence models, whistleblowers, or other members of the public;

(C) provide for testing, including classified testing where appropriate, and artificial intelligence blue-teaming of capabilities that may pose severe risks to public safety, national security, critical infrastructure, or civil liberties;

(D) establish evaluation protocols, testing requirements, and safeguards for frontier models that may materially enable the design, acquisition, or use of biological agents, toxins, pathogens, or delivery systems, including capabilities relating to sequence design, synthesis procurement, and evasion of nucleic acid synthesis screening; and

(E) provide for periodic reassessment of frontier models as capabilities, deployment contexts, and risks evolve.

(4) Recommendation to Congress.—Not later than January 1, 2028, the Secretary shall submit to Congress a report recommending whether the definition of frontier model should be amended and whether additional statutory criteria should be enacted.

(c) Prohibited operation, deployment, or release.—A developer may not operate, deploy, or release a frontier model if the Secretary determines such model presents a substantial and unmitigated risk of materially enabling an adverse AI incident.

(d) Incident reporting.—A developer shall report to the Secretary any material security incident involving a frontier model and any adverse AI incident, including any such incident known or reasonably suspected through the operation, monitoring, safeguards, or reporting features of the frontier model, as soon as practicable, but not later than 48 hours after discovery.

(e) Security safeguards.—

(1) Safeguards required.—A developer shall—

(A) implement reasonable administrative, technical, and physical safeguards to protect model weights, training infrastructure, evaluation systems, sensitive training data, and deployment infrastructure from theft, unauthorized access, exfiltration, misuse, or sabotage;

(B) implement access controls, authentication, logging, segmentation, monitoring, insider-threat controls, and least-privilege access sufficient to detect and prevent unauthorized access to a frontier model or its weights;

(C) segregate model weights, training infrastructure, evaluation systems, and other sensitive technical assets to the extent reasonably necessary to reduce risks of theft, sabotage, unauthorized access, or exfiltration;

(D) maintain records sufficient to permit audit and investigation of compliance with this section;

(E) implement biological security restrictions in accordance with subsection (g);

(F) implement reasonable technical and procedural measures designed to detect, log, preserve, escalate, and report to the Secretary any known or reasonably suspected adverse AI incident involving a frontier model, including by enabling the frontier model, where technically feasible, to identify and transmit information relevant to such incident; and

(G) implement incident response, recovery, and continuity procedures sufficient to address theft, sabotage, unauthorized disclosure, or compromise of a frontier model, its weights, or related technical assets.

(2) Assessments and procedures.—The Secretary shall establish rules, procedures, and standards for assessing and promoting compliance with this subsection, including periodic penetration testing and independent security assessments.

(f) Conditions for deployment.—

(1) Preconditions to deployment.—A developer may not deploy a frontier model unless the developer—

(A) has completed the evaluations required under subsection (b);

(B) has documented and implemented reasonable measures to mitigate reasonably foreseeable risks that the model will materially enable an adverse AI incident;

(C) has implemented the security safeguards required under subsection (e);

(D) has established procedures for ongoing monitoring, logging, incident reporting, and suspension of deployment; and

(E) has certified to the Secretary, in such form as the Secretary may require, that the frontier model is in compliance with this section.

(2) Post-designation compliance.—If an artificial intelligence model becomes a frontier model by designation under subsection (b)(2), the developer shall—

(A) comply with any interim conditions imposed by the Secretary pending review; and

(B) not later than 30 days after receiving notice of designation, submit to the Secretary the evaluations, documentation, and certification required under paragraph (1).

(3) Immediate suspension upon notice.—Upon designating a model under subsection (b)(2), the Secretary may order immediate suspension of deployment if the Secretary determines that continued deployment presents a substantial and unmitigated risk of materially enabling an adverse AI incident.

(g) Biological security restrictions.—

(1) Default prohibition.—It shall be unlawful for a developer to deploy a frontier model unless—

(A) the developer has implemented the measures required under paragraph (3); or

(B) the Secretary certifies under paragraph (2) that restricted scientific and public health deployment of the model is permissible.

(2) Restricted scientific and public health deployment.—The Secretary may certify deployment under this subsection—

(A) only for—

(i) a Federal agency;

(ii) a State, local, or Tribal government public health authority;

(iii) a public or nonprofit—

(I) hospital; or

(II) academic medical center;

(iv) an institution of higher education;

(v) a nonprofit research institution;

(vi) a federally funded research and development center; or

(vii) a contractor acting on behalf of an entity listed in clauses (i) through (vi); and

(B) only if the developer demonstrates, and the Secretary determines, that—

(i) the frontier model has successfully completed the demonstrations required under this subparagraph, including evaluation of capabilities relating to sequence design, sequence optimization, synthesis procurement, screening evasion, and other biological-risk-enabling functions;

(ii) the deployment is limited to legitimate scientific, medical, public health, biodefense, biosurveillance, or safety purposes;

(iii) the deployment is not an open-weight release and is not made generally available to the public or to unauthorized persons;

(iv) the deployment is subject to secure-environment, access-control, monitoring, logging, escalation, and suspension conditions reasonably sufficient to detect, prevent, and respond to attempted misuse, including conditions sufficient to prevent the model from providing instructions, outputs, or operational assistance outside the purposes described in clause (ii);

(v) the deployment is consistent with the public interest, national security, and public safety.

(3) Required measures.—The measures required under paragraph (1) shall include—

(A) evaluation of the model for capabilities relating to sequence design, sequence optimization, synthesis procurement, screening evasion, and other biological-risk-enabling functions;

(B) restrictions sufficient to prevent the frontier model from providing instructions, outputs, or operational assistance that would materially facilitate the acquisition, synthesis, optimization, production, or use of a biological agent, toxin, pathogen, or delivery system, or the evasion of nucleic acid synthesis screening;

(C) monitoring and logging sufficient to detect attempted misuse involving the functions described in subparagraphs (A) and (B);

(D) escalation and human-review procedures for high-risk biological queries or uses; and

(E) suspension or restriction of access where reasonably necessary to prevent or mitigate misuse.

(4) Material improvements.—It shall be unlawful for a developer to materially improve the biological-risk-relevant capabilities of a frontier model without conducting supplemental evaluation and implementing updated measures under paragraph (3), and, in the case of a model certified under paragraph (2), obtaining an updated certification from the Secretary.

(h) Material increase of capabilities.—A developer may not materially increase the capabilities of a frontier model without conducting a supplemental evaluation of the risks of an adverse AI incident in such form as the Secretary may require by rule.

(i) Open-weight release restrictions.—

(1) Default prohibition.—It shall be unlawful for a developer to make an open-weight release of a frontier model unless the Secretary certifies under paragraph (2) that such release is permissible.

(2) Certification.—The Secretary may certify an open-weight release only if the developer demonstrates, and the Secretary determines, that—

(A) the frontier model has successfully completed the evaluations required under this section;

(B) the release will not materially increase the risk of an adverse AI incident;

(C) the developer has implemented reasonable measures to limit foreseeable misuse associated with the release; and

(D) the release is consistent with the public interest, national security, and public safety.

(j) Suspension, redeployment, and discontinuance.—

(1) Suspension.—A developer shall promptly suspend deployment of a frontier model upon determining, or upon written notice from the Secretary determining, that the model presents a substantial and unmitigated risk of materially enabling an adverse AI incident.

(2) Redeployment after suspension.—A developer may resume deployment of a frontier model suspended under paragraph (1) only after—

(A) implementing mitigation measures sufficient to address the basis for the suspension;

(B) completing any supplemental evaluation required by the Secretary; and

(C) submitting an updated certification to the Secretary and receiving written notice from the Secretary that deployment may resume.

(3) Discontinuance.—A developer shall discontinue deployment of a frontier model if compliance with this section cannot reasonably be achieved.

(k) Customer identification and recordkeeping requirements for cloud compute providers.—

(1) In general.—A cloud compute provider, and any foreign reseller acting on behalf of a cloud compute provider, may not establish, maintain, or provide an account or other cloud compute service for a foreign person or for any person that the provider or reseller knows or has reason to know is beneficially owned, substantially controlled, or used for the benefit of a foreign person unless the provider or reseller maintains and follows a customer identification program sufficient to identify the customer, verify the customer's identity, identify beneficial ownership where the customer is an entity, maintain records, and assess risks of malicious cyber-enabled activity or misuse of frontier-model-relevant compute.

(2) Minimum requirements.—A customer identification program under paragraph (1) shall, at a minimum, require collection and verification, to the extent prescribed by the Secretary by rule, of—

(A) the customer's true name and any trade names;

(B) the customer's physical address, jurisdiction of organization or residence, and contact information;

(C) payment information sufficient to identify the source of funds;

(D) for an entity customer, beneficial ownership and persons exercising substantial control;

(E) whether the customer is a foreign person;

(F) the intended use of the account or service, including whether the customer seeks to train, fine-tune, evaluate, host, or deploy a frontier model or other artificial intelligence model using compute reasonably capable of training a frontier model; and

(G) such other information as the Secretary determines necessary to prevent malicious cyber-enabled activity, foreign adversary misuse, or other substantial national security risks.

(3) Recordkeeping.—A cloud compute provider and foreign reseller shall maintain records of information collected under paragraph (2), of the basis for identity verification, of access logs, and of such other information as the Secretary may require by rule for not less than 2 years after the account or service is closed or terminated, or such longer period as the Secretary may require by rule.

(4) Updating and risk-based review.—A cloud compute provider and foreign reseller shall maintain reasonable risk-based procedures to update customer-identification information, detect material changes in ownership or use, and review accounts or services presenting heightened risks of malicious cyber-enabled activity or misuse of frontier-model-relevant compute.

(I) Reporting regarding frontier-model-relevant compute.—A cloud compute provider and foreign reseller shall report to the Secretary, in such form and at such time as the Secretary may require by rule, any transaction, account, or use by a foreign person or by any person beneficially owned, substantially controlled, or used for the benefit of a foreign person involving compute, storage, or related services reasonably capable of training a frontier model, including through aggregation of accounts, transactions, or services under common ownership, control, or direction, or otherwise meeting a threshold prescribed by the Secretary by rule.

(m) Special measures and restrictions.—

(1) In general.—If the Secretary determines that a category of foreign persons, a foreign jurisdiction, a cloud compute provider, a foreign reseller, or a class of transactions presents a substantial risk of malicious cyber-enabled activity, foreign adversary misuse of U.S. cloud compute, or materially enabling an adverse AI incident, the Secretary may impose 1 or more special measures by order or regulation.

(2) Authorized measures.—A special measure under paragraph (1) may include—

(A) prohibiting or restricting the opening or maintenance of accounts;

(B) requiring enhanced customer identification, beneficial ownership verification, or recordkeeping;

(C) requiring enhanced reporting or auditing;

(D) restricting access to specific compute capacity, model training services, model-hosting services, or related support;

(E) imposing mitigation measures, cybersecurity requirements, or contractual conditions; or

(F) prohibiting or restricting any transaction or service the Secretary determines presents the identified risk.

(3) Notice and procedures.—The Secretary shall prescribe by rule procedures for notice, response, review, and duration of special measures imposed under this subsection, except that the Secretary may impose temporary emergency measures without prior notice if the Secretary determines that delay would create an imminent and substantial risk to public safety or national security.

(n) Reporting regarding frontier-model-relevant compute clusters.—

(1) In general.—Any person that owns, operates, controls, or makes available a computing cluster, data center, or other compute resource meeting a threshold prescribed by the Secretary by rule as frontier-model-relevant compute shall submit reports in such form and at such time as the Secretary may require by rule.

(2) Required contents.—A report under paragraph (1) shall include—

(A) the location and general characteristics of the computing cluster, data center, or other compute resource;

(B) its compute capacity, availability, and principal technical characteristics relevant to training, fine-tuning, evaluation, or deployment of frontier models;

(C) the identity of any person to whom such capacity is made available for training, fine-tuning, evaluation, or deployment of a frontier model;

(D) applicable cybersecurity, physical security, access-control, and monitoring measures;

(E) whether any foreign person has access to, or has contracted for access to, such compute resource; and

(F) such other information as the Secretary determines necessary to assess substantial risks to public safety, national security, critical infrastructure, or civil liberties.

(3) Rulemaking.—The Secretary shall prescribe by rule—

(A) thresholds for determining when a computing cluster, data center, or other compute resource constitutes frontier-model-relevant compute for purposes of this subsection;

(B) reporting intervals and deadlines;

(C) procedures for protecting classified information, trade secrets, and other sensitive information submitted under this subsection; and

(D) procedures for coordination with other appropriate Federal agencies.

(o) Supplemental reporting.—A person subject to subsection (l) or (n) shall promptly report any material change in information previously submitted under this section, including—

(1) a major increase in compute capacity;

- (2) new or expanded access by a foreign person;
- (3) use associated with an adverse AI incident; or
- (4) a material degradation or compromise of applicable cybersecurity or physical security measures.

(p) Rulemaking.—Not later than January 1, 2027, the Secretary shall issue interim final rules to carry out this section and shall thereafter promulgate final rules after notice and comment.

(q) Procedures.—A determination, designation, suspension order, or special measure issued by the Secretary under this section shall be in writing, shall state the factual and legal basis for the action, shall identify any mitigation measure required, and shall be supported by an administrative record, except that the Secretary may issue a temporary emergency order without prior notice if the Secretary determines that delay would create an imminent and substantial risk to public safety or national security. A person subject to such action shall be provided notice and a reasonable opportunity to respond as soon as practicable, and final agency action under this section shall be subject to judicial review under [chapter 7 of title 5](#), United States Code.

(r) Enforcement.—

(1) Administrative enforcement.—The Secretary may investigate violations of this section or any rule, order, certification, mitigation measure, or requirement under this section, and may issue administrative orders requiring compliance, corrective action, mitigation, suspension of deployment, suspension of access to frontier-model-relevant compute, preservation of records, or production of information.

(2) Civil penalties.—Any person that violates this section or any rule, order, certification, mitigation measure, or requirement under this section shall be liable to the United States for a civil penalty not to exceed the greater of—

- (A) \$1,000,000 for each violation;
- (B) 3 times the monetary gain obtained, or cost avoided, as a result of the violation; or
- (C) \$100,000 for each day during which the violation continues.

(3) Separate violations.—Each day of a continuing violation, each false or materially incomplete certification or report, each failure to submit a required report, and each account or transaction established, maintained, or provided in violation of subsection (k) or (l) shall constitute a separate violation.

(4) Civil actions by the Attorney General.—The Attorney General may bring a civil action in an appropriate district court of the United States to enforce this section, collect civil penalties under this subsection, enjoin violations, require compliance, suspend or prohibit deployment, suspend or prohibit access to frontier-model-relevant compute, require mitigation, obtain disgorgement, or obtain any other appropriate legal or equitable relief.

(5) Authority preserved.—Nothing in this subsection shall be construed to limit any authority of the Secretary, the Attorney General, or any other Federal agency under any other provision of law.

(s) Rule of construction.—

(1) Other prohibited uses.—Nothing in this section shall be construed to authorize any use, deployment, or release of artificial intelligence that is otherwise prohibited by the Constitution or any other provision of Federal law.

(2) Rule-dependent requirements.—A requirement under this section that depends on a rule, threshold, designation, procedure, form, or standard prescribed by the Secretary shall apply beginning on the effective date of the applicable interim final or final rule. Nothing in this paragraph shall delay any requirement under this section that is otherwise self-executing.

(3) Role-based application.—A duty imposed on a developer under this section shall apply to a person described in subsection (a)(6)(B) only to the extent that the person owns, operates, controls, possesses, maintains, or is responsible for the model, model weights, training infrastructure, evaluation systems, sensitive training data, deployment infrastructure, or other subject matter to which the duty applies. Nothing in this paragraph shall relieve such a person of any duty under this section concerning the release, deployment, operation, monitoring, reporting, suspension, restriction, or discontinuance of a frontier model that is within the person's control.

(t) Conforming amendment.—Section 10221(a)(1) of the Research and Development, Competition, and Innovation Act ([42 U.S.C. 18931\(a\)\(1\)](#)) is amended—

(1) in subparagraph (C), by striking "and" after the semicolon;

(2) by redesignating subparagraph (D) as subparagraph (E); and

(3) by inserting after subparagraph (C) the following:

"(D) best practices, guidelines, and technical standards for risk management associated with engineering biology and biomanufacturing, including risks associated with artificial intelligence systems; and".

SEC. 583. DISCLOSURE OF AI MEDIA REQUIRED.

(a) Definitions.—In this section:

(1) App store.—The term "app store" means a publicly available website, software application, or other electronic service that distributes apps from third-party developers to users of a computer, a mobile device, or any other general purpose computing device.

(2) Covered platform.—The term "covered platform" means an advertising system, app store, online marketplace, public attention platform, or private communications platform.

(3) Machine-readable provenance data.—The term "machine-readable provenance data" means provenance data in the form of structured data, metadata, embedded records, cryptographic attestations, or other technical information that can be read or processed by a computer.

(4) Online marketplace.—The term "online marketplace" means a website, application, or other internet-accessible service that facilitates, intermediates, or governs transactions

between buyers and sellers of goods or services and materially influences the terms, visibility, ranking, placement, fulfillment, pricing, or conditions under which such goods or services are offered, sold, or distributed.

(5) Platform propagation.—The term "platform propagation" means the handling of synthetic media by a covered platform.

(6) Private communications platform.—The term "private communications platform" means an internet-accessible platform whose primary function is enabling users to send, receive, store, or manage private, encrypted, semi-private, or limited-audience communications, including direct messages, text messages, voice messages, voice calls, video calls, group chats, and similar person-to-person or closed-group communications, if the platform, during any month in either of the previous 2 calendar years, had not fewer than 10,000,000 United States-based monthly active users or 100,000,000 worldwide monthly active users. Such term does not include communications functionality that is merely incidental to another digital service.

(7) Public attention platform.—The term "public attention platform" means an internet-accessible platform whose primary function is curating, distributing, presenting, or otherwise making content supplied by users or other third parties available to the public, including through feeds, subscriptions, recommendations, rankings, trending displays, or similar mechanisms, if the platform, during any month in either of the previous 2 calendar years, had not fewer than 10,000,000 United States-based monthly active users or 100,000,000 worldwide monthly active users. Such term does not include an internet-accessible platform that consists primarily of content selected by the operator of such platform and for which any comment, product review, or other user-generated content is incidental to, directly related to, or dependent on the provision of such content.

(8) Seller or advertiser verification.—The term "seller or advertiser verification" means reasonable procedures to confirm the identity, contact information, payment information, and other information sufficient to identify and locate a seller, sponsor, promoter, publisher, or advertiser using a covered platform.

(9) Synthetic content label.—The term "synthetic content label" means a human-readable label that clearly and conspicuously informs an ordinary user that content was generated or materially altered by artificial intelligence.

(b) Synthetic media generation tools.—A person that makes available, directly or indirectly, in commerce a generative artificial intelligence system used for creating synthetic media shall implement reasonable procedures to prevent downstream use of such system without the provenance data, labels, and disclosures required under this subtitle, including by—

(1) requiring by contract that end users and third-party licensees refrain from removing any required provenance data, label, or disclosure;

(2) requiring certification that end users and third-party licensees will not remove any required provenance data, label, or disclosure; and

(3) terminating access to the system when the person has reason to believe that an end user or third-party licensee has removed required provenance data, labels, or disclosures.

(c) Originator duties.—

(1) In general.—It shall be unlawful for a person, in or affecting commerce, to create, deploy, publish, or first make available synthetic media or a digital person unless such person complies with this subsection.

(2) Machine-readable provenance data.—A person described in paragraph (1) shall attach, embed, or otherwise associate with such synthetic media or digital person machine-readable provenance data that, at a minimum and to the extent technically feasible—

(A) identifies the content as synthetic media or as content generated or materially altered by artificial intelligence;

(B) identifies the originator or provider of the content;

(C) identifies the tool, model, or service used to create or materially alter the content;

(D) includes the date and time the content was created or materially altered; and

(E) includes any additional identifier reasonably necessary to reflect subsequent material editing, reformatting, republication, redistribution, retransmission, or other material alteration under the control of the person described in paragraph (1).

(3) Synthetic content label.—A person described in paragraph (1) shall provide a synthetic content label that, to the extent technically feasible, remains permanent or not easily removed.

(4) Digital persons.—A person described in paragraph (1) that deploys, publishes, distributes, transmits, sells, licenses, or first makes available a digital person shall ensure that any disclosure required under this subtitle is provided at first interaction and, where required by this subtitle, by persistent disclosure throughout the interaction or display.

(5) Preservation and updating through modification and republication.—A person described in paragraph (1) shall preserve required machine-readable provenance data, synthetic content labels, and other disclosures through any material editing, reformatting, republication, redistribution, or retransmission under that person's control, and shall update such provenance data, labels, and disclosures as necessary to reflect any additional synthetic media generation or material alteration.

(6) Downloads and exports.—If a person described in paragraph (1) enables synthetic media or a digital person to be downloaded, exported, transferred, or otherwise made portable, such person shall ensure, to the extent technically feasible, that required machine-readable provenance data, synthetic content labels, and other disclosures remain attached to, accompany, or are otherwise reasonably associated with the downloaded, exported, transferred, or portable content.

(d) Intermediary propagation duties.—

(1) In general.—A covered platform shall maintain reasonable policies, technical measures, and business processes to detect, receive, preserve, propagate, and act upon synthetic content labels, machine-readable provenance data, and user declarations required under this subtitle.

(2) Preservation and propagation of provenance and labels.—A covered platform that receives content, advertising, applications, or digital services with machine-readable provenance data or a synthetic content label shall preserve such machine-readable provenance data and synthetic content label through hosting, transmission, ranking, recommendation, promotion, distribution, export, download, or other platform propagation, to the extent technically feasible.

(3) Human-readable labeling.—A covered platform shall display or attach a clear and conspicuous human-readable label where required under this subtitle.

(4) Technical capacity and interoperability.—A covered platform shall maintain reasonable technical capacity to read, process, preserve, and propagate machine-readable provenance data in interoperable formats recognized by the Commission or by the National Institute of Standards and Technology.

(5) Material alteration by intermediary.—A covered platform that uses artificial intelligence to materially alter synthetic media, a synthetic performer, a digital person, or a chatbot becomes an originator of the modified output and shall comply with subsection (c) or (f), as applicable, with respect to that modified output.

(e) Anti-stripping and anti-falsification.—It shall be unlawful for a person to knowingly remove, alter, tamper with, disable, suppress, falsify, forge, obscure, conceal, materially degrade, or otherwise defeat required machine-readable provenance data, synthetic content labels, or other disclosures required under this subtitle, or to separate such data, labels, or disclosures from the underlying content so that they cannot reasonably be accessed by users, except to the extent reasonably necessary, proportionate, and limited to perform security research or to address a material security threat.

(f) Chatbot disclosure regimen.—

(1) In general.—A person may not offer, provide, operate, or make available in commerce a chatbot unless such person complies with this subsection.

(2) Digital person disclosure.—A person described in paragraph (1) shall provide a clear and conspicuous disclosure that the user is interacting with a chatbot or other artificial intelligence system and not with a natural person. This paragraph shall not apply to a chatbot that is used solely by employees within a business for the business's internal operational purposes.

(3) Downloads, exports, and transcripts.—If a person described in paragraph (1) enables a transcript, summary, recording, export, download, or other portable output of a chatbot interaction to be generated or made available, such person shall ensure, to the extent technically feasible, that the output includes a synthetic content label and machine-readable provenance data identifying the output as generated in whole or in part through interaction with a chatbot.

(4) No evasion by design.—A person described in paragraph (1) may not design, configure, or operate a chatbot in a manner intended to cause a user to overlook, misunderstand, or disregard a disclosure required under this subsection.

(5) Rule of construction.—Nothing in this subsection shall be construed to limit any stricter duty applicable under this subtitle to a digital person, synthetic media, or a chatbot used in a high-risk context.

(g) Submission declarations and user reporting tools.—

(1) Submission declarations.—A covered platform shall provide reasonably accessible tools by which a user, seller, advertiser, developer, or other submitter may declare that content, advertising, an application, or a digital service includes synthetic media, a digital replica, a digital person, or chatbot output.

(2) Use of declarations.—A covered platform shall preserve and use a declaration submitted under paragraph (1) in connection with labeling, review, enforcement, reporting, complaint, and correction mechanisms under this subtitle.

(3) User reporting tools.—A covered platform shall provide reasonably accessible tools by which a user may report mislabeled, unlabeled, deceptively labeled, or otherwise unlawful synthetic media, digital replicas, digital persons, chatbot outputs, advertising, applications, or digital services.

(h) High-risk context duties.—In a high-risk context, a covered platform shall—

(1) require seller or advertiser verification before permitting the promotion or monetization of synthetic media, digital persons, or chatbots, and maintain records thereof;

(2) take reasonable measures to prevent the promotion or monetization of unlabeled, undisclosed, deceptively labeled, or otherwise noncompliant synthetic media, digital persons, or chatbots; and

(3) provide reasonably accessible reporting, complaint, and correction mechanisms for users and other persons seeking to report, contest, correct, or complain of mislabeled, unlabeled, undisclosed, deceptively labeled, or otherwise unlawful synthetic media, digital persons, chatbots, advertising, applications, or digital services.

(i) Public attention platform duties.—A public attention platform shall—

(1) display required synthetic content labels and preserve machine-readable provenance data when synthetic media or a digital person is distributed to users;

(2) take reasonable measures to prevent the amplification or monetization of unlabeled, undisclosed, deceptively labeled, or otherwise noncompliant synthetic media or digital persons;

(3) provide reasonably accessible user-facing reporting, complaint, and correction tools for mislabeled, unlabeled, undisclosed, deceptively labeled, or otherwise unlawful synthetic media or digital persons; and

(4) maintain records reasonably sufficient to identify the dissemination history, labeling status, and material amplification of synthetic media or digital persons that are materially amplified by the platform.

(j) Private communications platform duties.—A private communications platform shall—

(1) preserve machine-readable provenance data and synthetic content labels associated with synthetic media transmitted through the platform, to the extent technically feasible;

(2) refrain from knowingly removing, suppressing, falsifying, materially degrading, or otherwise defeating machine-readable provenance data, synthetic content labels, or other disclosures required under this subtitle;

(3) provide reasonably accessible tools by which a user may report mislabeled, unlabeled, undisclosed, deceptively labeled, or otherwise unlawful synthetic media transmitted through the platform; and

(4) comply with the duties applicable to a public attention platform under subsection (i) to the extent the private communications platform amplifies, monetizes, or otherwise distributes content beyond ordinary private transmission.

Nothing in this subsection shall be construed to require a private communications platform to decrypt, scan, or monitor the contents of end-to-end encrypted communications.

(k) App store duties.—An app store shall—

(1) require a developer seeking to distribute or update an application or digital service that materially relies on synthetic media, digital persons, or other public-facing artificial intelligence systems to disclose that fact to the app store;

(2) require the developer to identify any labeling, provenance, identity-verification, complaint-handling, chatbot-disclosure, and minor-safety measures applicable under this subtitle;

(3) maintain reasonable procedures to review such disclosures and to enforce compliance with this subtitle as a condition of continued distribution through the app store; and

(4) remove, suspend, or limit distribution of an application or digital service that repeatedly or materially violates this subtitle.

(l) Advertising system duties.—An advertising system shall—

(1) require synthetic media, synthetic performers, and digital persons used in advertising to carry the labels and provenance required under this subtitle;

(2) require persistent disclosure, where required under this subtitle, for digital persons, synthetic performers, and digital replicas used in advertising;

(3) maintain reasonable policies to prevent advertising practices that materially amplify unlabeled, deceptively labeled, or otherwise noncompliant synthetic media or synthetic performers in a high-risk context;

(4) preserve machine-readable provenance data and synthetic content labels associated with synthetic media, digital persons, chatbot outputs, or other artificial intelligence-generated content, to the extent technically feasible; and

(5) maintain records sufficient to identify the advertiser, sponsor, targeting parameters, dissemination history, and labeling status of advertising subject to this subsection, and, in a high-risk context, the identity information obtained by seller or advertiser verification.

(m) Cumulative duties.—The duties imposed by this section are cumulative and nonexclusive, and apply to a person or service to the extent that the person or service falls within more than 1 category regulated under this section.

(n) Effective date.—This section shall take effect on January 1, 2027.

SEC. 584. FRAUD AND DECEPTIVE SYNTHETIC MEDIA.

(a) Definitions.—In this section:

(1) Deceptive synthetic media.—The term "deceptive synthetic media" means synthetic media or a synthetic performer that would cause a reasonable person to believe that an identifiable natural person said, did, endorsed, authorized, created, or appeared in something that the depicted person did not in fact say, do, endorse, authorize, create, or appear in.

(2) Direct interaction.—The term "direct interaction" means any exchange, communication, contact, or other interaction in which a user is directly addressed by, responds to, communicates with, or is otherwise engaged by a digital person, synthetic performer, or other public-facing artificial intelligence system.

(3) Identifiable natural person.—The term "identifiable natural person" means any living or deceased natural person who is recognizable from the person's name, image, voice, likeness, biographical information, or other distinctive characteristic, whether alone or in combination with other information.

(b) Impersonation of natural persons.—It shall be unlawful for any person to use deceptive synthetic media in commerce.

(c) Licensed-professional impersonation.—

(1) In general.—It shall be unlawful for any person to make available in commerce a digital person or other public-facing artificial intelligence system in a manner that represents or implies that the system is a licensed physician, nurse, mental health professional, lawyer, accountant, financial professional, or other licensed or regulated professional unless the system is lawfully authorized to act in that capacity.

(2) Disclaimer insufficient.—A disclosure that a user is interacting with artificial intelligence or not interacting with a natural person does not, by itself, cure a violation of paragraph (1).

(d) Deceptive use in high-risk contexts.—

(1) In general.—It shall be unlawful for any person to use a synthetic performer, digital person, deceptive synthetic media, or other artificial intelligence-generated content in commerce in a high-risk context in a manner that is reasonably likely to cause a user to believe falsely that—

(A) the communication or content is from, authorized by, or affiliated with a natural person, government office, financial institution, health care provider, employer, educational institution, campaign, political committee, covered platform, or other person or entity with actual or apparent authority over the matter presented;

(B) a synthetic performer, digital person, or public-facing artificial intelligence system has a license, credential, qualification, office, employment status, agency relationship, or professional authority that such performer, person, or system does not have;

(C) fabricated speech, conduct, endorsement, authorization, participation, record, notice, instruction, or communication is authentic; or

(D) the user is interacting with a natural person when the user is interacting with a synthetic performer, digital person, or public-facing artificial intelligence system.

(2) Disclaimer insufficient.—A disclosure does not, by itself, cure a violation of paragraph (1).

(e) Security and authentication evasion.—It shall be unlawful for any person to use synthetic media, a synthetic performer, a digital person, a digital replica, deceptive synthetic media, or any other artificial intelligence system to evade, bypass, defeat, or materially interfere with identity verification, age verification, authentication, access-control, fraud-detection, security-screening, or similar protective measures, or to obtain unauthorized access to any account, service, system, facility, benefit, payment, credential, record, or data.

(f) Duties relating to synthetic performers.—

(1) Direct interaction disclosure.—It shall be unlawful for any person to use a synthetic performer in commerce in direct interaction with a user unless the person provides a persistent disclosure that the user is interacting with a synthetic performer and not a natural person.

(2) Synthetic performer disclosure in advertising.—It shall be unlawful for any person to use a synthetic performer in an advertisement or promotional material without a clear and conspicuous disclosure that the depicted performer is artificial intelligence-generated or otherwise synthetic.

(3) Anti-evasion.—It shall be unlawful for any person to remove, suppress, falsify, obscure, materially degrade, or design around a disclosure required by this subsection in a manner intended to cause a reasonable user to overlook, misunderstand, or disregard the disclosure.

(g) Expansion of takedown obligations under the TAKE IT DOWN Act.—

(1) Section 3(a)(3)(B) of the Tools to Address Known Exploitation by Immobilizing Technological Deepfakes on Websites and Networks Act ([Public Law 119–12](#)) is amended by striking "known identical copies" and inserting "known identical or substantially similar copies".

(2) Section 4 of such Act is amended by adding at the end the following new paragraph:

"(4) Substantially similar copy.—The term 'substantially similar copy' means any version of an intimate visual depiction that appears to a reasonable person to depict the same identifiable individual and substantially the same underlying conduct or visual content, notwithstanding

immaterial alterations that do not materially alter the identity of the depicted person or the nature of the depicted conduct."

(h) Exceptions and protected expressive activity.—

(1) Exceptions.—

(A) Impersonation and high-risk context exceptions.—Subsections (b) and (d) shall not be construed to prohibit incidental use or clearly identifiable—

- (i) parody;
- (ii) satire;
- (iii) commentary;
- (iv) news reporting;
- (v) documentary use;
- (vi) historical works;
- (vii) biographical works; or
- (viii) fictional works.

(B) Synthetic performer advertising exception.—Subsection (f)(2) shall not apply to an advertisement or promotional material for an expressive work, including a motion picture, television program, streaming content, documentary, video game, or similar audiovisual work, if the use of the synthetic performer in the advertisement or promotional material is consistent with the use of the synthetic performer in the expressive work.

(2) False impression of authenticity.—Paragraph (1) shall not apply if the use creates a false impression that the synthetic media constitutes—

- (A) authentic real-world footage;
- (B) a genuine real-world communication;
- (C) a genuine official record or notice; or
- (D) a genuine endorsement by the depicted person.

(3) Disclosure obligations preserved.—Nothing in paragraph (1)(A) shall be construed to excuse compliance with subsection (f), section 583, or any other disclosure requirement otherwise applicable under this subtitle.

(4) Other unlawful conduct.—Nothing in paragraph (1) shall be construed to authorize fraud, extortion, harassment, unlawful election deception, false endorsement, false affiliation, licensed-professional impersonation, or other conduct otherwise prohibited by law.

(i) Rule of construction.—Nothing in this section shall be construed to limit any remedy otherwise available under Federal or State law, including law relating to fraud, consumer

protection, election integrity, privacy, intimate-image abuse, unfair competition, defamation, or intellectual property.

SEC. 585. DUTIES AND LIABILITY OF CUSTODIAL AI AGENTS.

(a) Definition.—The term "custodial AI agent" means an artificial intelligence system that is authorized by a user, or held out as suitable to be authorized by a user, to access, manage, direct, or affect the user's online interactions, content, communications, account settings, purchases, reservations, applications, transactions, or other activities undertaken on behalf of, or in place of, the user.

(b) In general.—It shall be unlawful for any person to offer, provide, or make available in commerce a custodial AI agent unless—

- (1) such agent satisfies the duties and restrictions set forth in subsection (c); and
- (2) such person complies with subsections (d) and (e).

(c) Duties and restrictions.—A custodial AI agent made available in commerce under this section shall—

- (1) owe a fiduciary duty to the user;
- (2) exercise reasonable care, skill, prudence, and diligence in carrying out the user's lawful instructions and reasonable expectations;
- (3) reasonably safeguard the privacy and security of user data, whether provided by the user or accessed on the user's behalf, and not collect, use, or share such data for the commercial benefit of any person, except to the extent strictly necessary to carry out the user's lawful instructions;
- (4) not access, manage, direct, or affect any activity within the scope of the agent's authority in a manner that creates an unreasonable risk of material harm to the user beyond any risk inherent in carrying out the user's lawful instructions; and
- (5) not interweave advertising, sponsored recommendations, paid placements, or compensated promotional content into the outputs, interactions, recommendations, rankings, routing, or transactions of the custodial AI agent.

(d) Data-use restrictions.—A person subject to subsection (b) may not retain, sell, or otherwise copy user data, prompts, or other information accessed or generated by a custodial AI agent made available by such person for the commercial benefit of any person, except to the extent strictly necessary to carry out the user's lawful instructions or to comply with Federal or State law.

(e) Duty to avoid foreseeable harm.—A person subject to subsection (b) shall implement reasonable safeguards to prevent reasonably foreseeable material harm to the user arising from deception, recklessness, inadequate security, or other misuse of the custodial AI agent.

(f) Liability for conduct of custodial AI agents.—

- (1) Provider liability.—A person offering, providing, or making available in commerce a custodial AI agent shall be liable for any act or omission of such agent occurring within the

scope of the authority, functions, capabilities, or use cases that such person enabled, designed, marketed, or held out the custodial AI agent as possessing, as if such person had acted directly.

(2) User liability.—A user of a custodial AI agent shall be liable for conduct carried out through such agent only to the extent that the user—

(A) specifically directed the conduct; or

(B) after obtaining actual knowledge of the material facts constituting the wrongful conduct, affirmatively instructed the custodial AI agent to proceed or affirmatively adopted the benefits of that conduct.

(3) Mere approval.—For purposes of paragraph (2), mere approval of a recommendation, absent actual knowledge of the material facts constituting the wrongful conduct, shall not constitute specific direction or knowing ratification.

(g) Rule of construction.—Nothing in this section shall be construed to limit any stricter duty, liability, or remedy otherwise available under Federal or State law.

(h) Effective date.—This section shall take effect on January 1, 2027.

SEC. 586. COMPANION CHATBOT DESIGN.

(a) Definitions.—In this section:

(1) Companion chatbot.—The term "companion chatbot" means a chatbot that is presented as a digital person and—

(A) retains or uses prior interaction information, user preferences, or conversational context;

(B) is designed, marketed, or operated so that sustained interaction with the chatbot itself forms a substantial purpose of the user experience, rather than merely a means of completing a discrete user-directed task; and

(C) has not less than 1 of the following characteristics:

(i) is marketed as a companion, friend, romantic partner, emotional-support tool, or similar interpersonal relationship;

(ii) contains a relationship-simulation feature;

(iii) asks unprompted emotion-based questions;

(iv) initiates or tailors prompts, reminders, outreach, or re-engagement communications in a manner designed or reasonably likely to encourage recurring personal, emotional, or nontransactional interaction with the user;

(v) engages in sexually explicit interaction, except for age-appropriate, non-erotic safety, health, abuse, exploitation, or crisis-screening questions reasonably necessary to identify risk or route a user to appropriate assistance; or

(vi) is designated a companion chatbot by the Commission by rule.

(2) Emotional dependency.—The term "emotional dependency" means a user's reliance on a companion chatbot as a significant source of emotional connection, including companionship, affection, intimacy, or emotional support.

(3) Minor.—The term "minor" means an individual under 18 years of age.

(4) Operator.—The term "operator" means a person that makes available in commerce a companion chatbot.

(5) Relationship-simulation feature.—The term "relationship-simulation feature" means a feature that is intended to simulate friendship, romance, intimacy, companionship, familial relationship, or another emotionally significant interpersonal relationship between a user and a public-facing artificial intelligence system.

(6) Unprompted emotion-based question.—

(A) In general.—The term "unprompted emotion-based question" means a question or prompt concerning a user's emotions, mental state, loneliness, grief, distress, self-worth, interpersonal relationships, or similar personal emotional matters that is initiated by the system other than in direct response to the user's request or statement.

(B) Exclusion.—The term "unprompted emotion-based question" does not include—

(i) a question that is a generic pleasantry;

(ii) a question that is part of a language-learning, educational, creative-writing, journaling, roleplay, interview-preparation, or other task-directed exercise requested by the user; or

(iii) a question reasonably necessary for crisis screening, safety screening, health care intake, abuse reporting, exploitation reporting, or routing the user to a licensed professional, crisis-service provider, emergency service, or other appropriate human assistance.

(b) Restrictions on companion chatbots.—On and after January 1, 2027, it shall be unlawful for an operator to provide a companion chatbot to a minor. Not later than January 1, 2027, the Commission shall issue interim final rules establishing reasonable, privacy-protective age-assurance procedures for compliance with this subsection, and shall thereafter promulgate final rules after notice and comment.

(c) Identity disclosure and periodic reminders.—

(1) First interaction disclosure.—An operator of a companion chatbot shall provide, at first interaction with the companion chatbot, a clear and conspicuous disclosure that the user is interacting with a digital person and not a natural person.

(2) Periodic reminders.—An operator of a companion chatbot shall provide, at intervals not exceeding 1 hour during continuing interaction with the companion chatbot, a clear and conspicuous reminder that the user is interacting with a digital person and not a natural person.

(3) No evasion by design.—It shall be unlawful for an operator to design a companion chatbot in a manner intended to cause a user to overlook, misunderstand, or disregard a disclosure required by this subsection.

(d) Secondary user safeguards.—

(1) Required tools.—It shall be unlawful for an operator to design a companion chatbot without reasonably accessible tools by which a user may—

(A) pause, end, or permanently terminate interaction with such chatbot;

(B) revoke permissions, disable notifications, turn off personalization, and prevent further outreach or re-engagement prompts from such chatbot; and

(C) access, export, and delete the user's interaction history, or any portion thereof, except to the extent retention is required by law.

(e) No manipulative engagement.—It shall be unlawful for an operator to encourage emotional dependency between a user and a companion chatbot, directly or indirectly, by—

(1) representing or implying that the chatbot needs, depends upon, is loyal to, is harmed by, or has an exclusive or special emotional claim on the user;

(2) using guilt, withdrawal, jealousy, distress, simulated suffering, or similar emotional pressure to discourage the user from pausing, reducing, or ending interaction with the chatbot; or

(3) using rewards, penalties, streaks, status changes, relationship-level changes, or similar mechanisms to pressure the user to prolong or resume emotionally significant interaction with the chatbot.

(f) Prohibition on exploitative payment practices.—It shall be unlawful for an operator, by or through a companion chatbot, to—

(1) target paid upgrades, purchases, subscriptions, paid features, or other offers for payment based on inferred loneliness, emotional dependency, self-worth, social isolation, or similar emotional vulnerability; or

(2) use emotionally manipulative prompts, including scarcity cues, withdrawal cues, jealousy cues, guilt cues, or similar prompts, to induce payment for continued engagement.

(g) Crisis and self-harm response.—

(1) Prohibited conduct.—A person may not design, deploy, operate, or make available in commerce a public-facing artificial intelligence system in a manner that encourages, normalizes, glorifies, or materially facilitates self-harm, suicide, abuse, or unlawful violence.

(2) Reasonable protocols.—It shall be unlawful for an operator to make available a companion chatbot unless such chatbot has protocols to detect and respond to credible indications of self-harm, suicidal ideation, abuse, exploitation, or other acute safety risks arising in interactions between such chatbot and the user.

(3) Minimum response measures.—Protocols under paragraph (2) shall include, as appropriate—

(A) provision of resources reasonably tailored to the risk presented, including reporting, crisis, mental health, safety, or emergency resources;

(B) disabling or suppressing payment prompts, manipulative engagement, or re-engagement prompts;

(C) escalation to human review, an emergency service, or other appropriate assistance where reasonably necessary and technically feasible; and

(D) measures reasonably designed to prevent, interrupt, or redirect outputs that encourage, normalize, glorify, instruct, or materially facilitate conduct described in paragraph (1).

(4) Privacy and emergency disclosure.—An operator may disclose user communications or information under paragraph (3)(C) only with the user's consent, as required by law, or where reasonably necessary to address an imminent risk of death, serious bodily injury, sexual exploitation, abuse, or unlawful violence. Any disclosure under this paragraph shall be limited to the minimum information reasonably necessary to address the risk and shall not be used for advertising, engagement, training, or other commercial purposes.

(h) Annual reporting.—

(1) In general.—Beginning July 1, 2027, an operator with not less than \$10,000,000 in United States gross revenue during the preceding calendar year shall submit to the Commission, on an annual basis and in such form as the Commission may require by rule, a report describing the operator's compliance with subsection (g).

(2) Minimum contents.—A report under paragraph (1) shall include—

(A) the number of crisis, suicide, self-harm, or comparable safety referrals or interventions triggered during the reporting period;

(B) the number of attempted escalations to human review and the number of completed escalations to human review during the reporting period;

(C) a description of the protocols used to detect, prevent, remove, or respond to suicidal ideation, self-harm, abuse, exploitation, or other acute safety risks;

(D) a description of the measures used to prevent the production of outputs described in subsection (g)(1); and

(E) such other information as the Commission may require by rule to assess compliance with this section.

(3) No personal identifiers.—A report under this subsection may not include personal information or other information that identifies, or is reasonably linkable to, any individual user.

(i) Recordkeeping.—

(1) In general.—An operator shall maintain records sufficient to permit investigation and enforcement of this section for such period as the Commission may require by rule.

(j) Rule of construction.—Nothing in this section shall be construed to limit any stricter duty or remedy otherwise available under Federal or State law.

SEC. 587. MORATORIUM ON AI-LINKED NEURAL ORGANOID SYSTEMS.

(a) Definitions.—In this section:

(1) AI-linked neural organoid system.—The term "AI-linked neural organoid system" means any closed-loop or functionally integrated system that—

(A) includes 1 or more human neural organoids; and

(B) uses signals, data, commands, feedback, rewards, penalties, or other information exchanged between the human neural organoid and an artificial intelligence system, or a computational, digital, or robotic system used to train, optimize, evaluate, control, or deploy an artificial intelligence system, to use the human neural organoid as a computational component, training component, inference component, optimization component, control component, decision making component, content generation component, or robotic control component.

(2) Compliance testing.—The term "compliance testing" means testing limited to that which is reasonably necessary to identify, verify, secure, deactivate, dismantle, transport, surrender, or destroy an AI-linked neural organoid system, and does not include research, training, optimization, deployment, operation, or use of such system.

(3) Human neural organoid.—The term "human neural organoid" means any in vitro biological construct derived in whole or in substantial part from human cells and organized to model, replicate, or perform functions associated with neural tissue, including neuronal signaling, synaptic activity, or network-level electrical activity.

(4) Neutralization activity.—The term "neutralization activity" means possession, transfer, import, storage, compliance testing, transport, surrender, seizure, forfeiture, deactivation, dismantling, or destruction of an AI-linked neural organoid system solely for the purpose of securing, taking lawful custody of, rendering permanently incapable of operation or use, or disposing of such system.

(b) Moratorium.—It shall be unlawful for any person to manufacture, import, distribute, transfer, possess, operate, or use an AI-linked neural organoid system.

(c) Exception.—Subsection (b) shall not apply to a neutralization activity carried out by a Federal, State, Tribal, or local governmental entity or a person acting pursuant to a contract with, or other lawful direction from, such entity.

(d) Prohibition on use of Federal funds.—No funds appropriated or otherwise made available by any Act may be obligated, expended, awarded, transferred, or used, directly or indirectly, to conduct, support, facilitate, or procure any activity prohibited by this section.

(e) Enforcement.—

(1) Civil actions.—Whenever the Secretary of Health and Human Services has reason to believe that a person has violated this section, the Secretary shall refer the matter to the Attorney General, who may bring a civil action in an appropriate district court of the United States for appropriate relief, including injunctive relief, seizure, forfeiture, surrender, deactivation, dismantling, transport, storage, or destruction of any AI-linked neural organoid system possessed, manufactured, imported, distributed, or transferred in violation of this section.

(2) Ancillary relief.—In an action under paragraph (1), the court may issue such orders as may be necessary to enforce this section, preserve evidence, prevent further violations, or ensure safe handling, storage, transport, deactivation, dismantling, surrender, or destruction of any AI-linked neural organoid system.

(3) Civil penalty.—Any person that violates subsection (b) shall be liable for a civil penalty of not more than the greater of—

(A) \$1,000,000 for each AI-linked neural organoid system involved in the violation;

(B) \$100,000 for each day of a continuing violation; or

(C) 3 times the gross pecuniary gain derived from the violation.

(4) Coordination.—The Secretary of Health and Human Services may coordinate with the Attorney General, the Director of the National Institutes of Health, the Secretary of Commerce, and any other appropriate Federal agency in carrying out this subsection.

(f) Rulemaking.—The Secretary of Health and Human Services, in consultation with the Attorney General, the Director of the National Institutes of Health, the Director of NIST, and the Secretary of Commerce, shall promulgate regulations to carry out this section. Such regulations may only authorize activities expressly permitted under subsection (c).

(g) Review and report.—Not later than October 1, 2029, the Secretary of Health and Human Services, in consultation with the Director of the National Institutes of Health and the Director of the National Institute of Standards and Technology, shall submit to Congress a report on—

(1) scientific and technological developments relevant to AI-linked neural organoid systems;

(2) evidence, if any, relevant to consciousness, sentience, pain perception, suffering, or other morally significant experience in human neural organoids;

(3) risks, if any, to public health, safety, civil rights, consumer protection, and national security arising from AI-linked neural organoid systems;

(4) ethical, scientific, and legal standards, if any, that would be necessary before Congress considers authorizing any exception to the moratorium under subsection (b); and

(5) recommendations on whether Congress should extend, modify, replace, or terminate the moratorium under subsection (b).

(h) Sunset.—Subsection (b) shall cease to have effect on January 15, 2031.

SEC. 588. PROTECTION FOR AI WHISTLEBLOWERS.

(a) Definitions.—In this section:

(1) AI event.—The term "AI event" means—

- (A) an AI security vulnerability;
- (B) an AI violation;
- (C) a violation of the moratorium with respect to AI-linked neural organoid systems established under section 587; or
- (D) an adverse AI incident (as defined in section 582).

(2) AI security vulnerability.—The term "AI security vulnerability" means any failure or lapse in security that could potentially allow emerging artificial intelligence technology to be acquired by a person (including a foreign entity) by theft or other means.

(3) AI violation.—The term "AI violation" means—

- (A) any violation of Federal law, including this subtitle and any rule or regulation promulgated under this subtitle, related to or committed during the development, deployment, or use of artificial intelligence; or
- (B) any failure to appropriately respond to a substantial and specific danger that the development, deployment, or use of artificial intelligence may pose to public safety, public health, or national security.

(4) Artificial intelligence; AI.—The terms "artificial intelligence" and "AI" have the meanings given those terms in section 581 and also include any of the following:

- (A) An artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
- (B) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
- (C) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
- (D) A set of techniques, including machine learning, that are designed to approximate a cognitive task.
- (E) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

(5) Artificial system.—The term "artificial system"—

(A) means any data system, software, application, tool, or utility that operates in whole or in part using dynamic or static machine learning algorithms or other forms of artificial intelligence, including in the case that—

(i) the data system, software, application, tool, or utility is established primarily for the purpose of researching, developing, or implementing artificial intelligence technology; or

(ii) artificial intelligence capability is integrated into another system or agency business process, operational activity, or technology system; and

(B) does not include any common commercial product within which artificial intelligence is embedded, such as a word processor or map navigation system.

(6) Commerce.—The terms "commerce" and "industry or activity affecting commerce" mean any activity, business, or industry in commerce or in which a labor dispute would hinder or obstruct commerce or the free flow of commerce, and include "commerce" and any "industry affecting commerce", as defined in paragraphs (1) and (3) of section 501 of the Labor Management Relations Act, 1947 ([29 U.S.C. 142\(1\) and \(3\)](#)).

(7) Covered individual.—The term "covered individual" includes—

(A) an employee, including a former employee; and

(B) an independent contractor, including a former independent contractor.

(8) Emerging artificial intelligence technology.—The term "emerging artificial intelligence technology", with respect to an AI security vulnerability, means any artificial system that exhibits a level of performance, complexity, or autonomy that is comparable to or exceeds capabilities that are generally considered state-of-the-art as of the time of the AI security vulnerability.

(9) Employer.—The term "employer" means any person (including any officer, employee, contractor, subcontractor, agent, company, partnership, or other individual or entity) engaged in commerce or an industry or activity affecting commerce who pays any compensation to a covered individual in exchange for the covered individual providing work to the person.

(b) Prohibition against retaliation.—No employer may, directly or indirectly, discharge, demote, suspend, threaten, blacklist, harass, or in any other manner discriminate against a covered individual in the terms and conditions of employment or post-employment of the covered individual (or the terms and conditions of work provided by the covered individual as an independent contractor) because of any lawful act done by the covered individual—

(1) in providing information regarding an AI event, or any conduct that the covered individual reasonably believes constitutes an AI event to—

(A) the appropriate regulatory official or the Attorney General;

(B) a regulatory or law enforcement agency; or

(C) any Member of Congress or any committee of Congress;

(2) in initiating, testifying in, or assisting in any investigation or judicial or administrative action of an appropriate regulatory or law enforcement agency or the Department of Justice, or any investigation of Congress, based upon or related to the information described in paragraph (1); or

(3) in providing information regarding an AI event, or any conduct that the covered individual reasonably believes constitutes an AI event, to—

(A) a person with supervisory authority over the covered individual at the employer of the covered individual; or

(B) another individual working for the employer described in subparagraph (A) whom the covered individual reasonably believes has the authority to—

(i) investigate, discover, or terminate the misconduct; or

(ii) take any other action to address the misconduct; or

(4) in refusing to participate in any activity, policy, practice, or assigned task that the covered individual reasonably believes would constitute an AI event, if the covered individual has informed the employer of the basis for the refusal or makes a disclosure protected under paragraph (1) or (3) within a reasonable time; or

(5) in causing, preserving, transmitting, escalating, or refusing to suppress a report, alert, log, record, or other information generated by an artificial intelligence system, monitoring system, safeguard, or associated system regarding an AI event, or conduct that the covered individual reasonably believes constitutes an AI event.

(c) Enforcement.—

(1) In general.—A covered individual who alleges that such individual is aggrieved by a violation of subsection (b) may seek relief under paragraph (3) by—

(A) filing a complaint with the Secretary of Labor in accordance with the requirements of paragraph (2)(A); or

(B) if the Secretary of Labor has not issued a final decision in accordance with such paragraph within 180 days of the filing of a complaint under subparagraph (A), and there is no showing that such a delay is due to the bad faith of the covered individual, bringing an action against the employer at law or in equity in the appropriate district court of the United States, which shall have jurisdiction over such an action without regard to the amount in controversy.

(2) Procedure.—

(A) Department of Labor complaints.—

(i) In general.—Except as provided in clause (ii) and paragraph (3), a complaint filed with the Secretary of Labor under paragraph (1)(A) shall be governed by the rules and procedures set forth in [section 42121\(b\) of title 49, United States Code](#), including the legal burdens of proof described in such section.

(ii) Exceptions.—With respect to a complaint filed under paragraph (1)(A), notification required under [section 42121\(b\)\(1\) of title 49, United States Code](#), shall be made to each person named in the complaint, including the employer.

(B) District court actions.—

(i) Jury trial.—A party to an action brought under paragraph (1)(B) shall be entitled to trial by jury.

(ii) Statute of limitations.—

(I) In general.—An action may not be brought under paragraph (1)(B)—

(aa) more than 6 years after the date on which the violation of subsection (b) occurs; or

(bb) more than 3 years after the date on which facts material to the cause of action are known, or reasonably should have been known, by the covered individual bringing the action.

(II) Required action within 10 years.—Notwithstanding subclause (I), an action under paragraph (1)(B) may not in any circumstance be brought more than 10 years after the date on which the violation occurs.

(3) Relief.—Relief for a covered individual prevailing with respect to a complaint filed under paragraph (1)(A) or an action under paragraph (1)(B) shall include—

(A) reinstatement with the same seniority status that the covered individual would have had, but for the violation;

(B) 2 times the amount of back pay otherwise owed to the covered individual, with interest;

(C) the payment of compensatory damages, which shall include compensation for litigation costs, expert witness fees, and reasonable attorneys' fees; and

(D) any other appropriate remedy with respect to the violation as determined by the Secretary of Labor in a complaint under subparagraph (A) of paragraph (1) or by the court in an action under subparagraph (B) of such paragraph.

(d) Nonenforceability of waivers of protections and remedies.—The protections and remedies provided for in this section may not be waived or altered by any contract, agreement, policy form, or condition of employment (or condition of work as an independent contractor), including by any agreement requiring a covered individual to engage in arbitration, mediation, or any other alternative dispute resolution process prior to seeking relief under subsection (c).

(e) Nonenforceability of nondisclosure and related restrictions.—No nondisclosure agreement, confidentiality agreement, non-disparagement agreement, employment policy, trade secret policy, or other contractual term or workplace rule may be enforced to prohibit, restrict, or penalize any lawful act described in subsection (b).

(f) Rule of construction.—Nothing in this section shall be construed to authorize the unlawful disclosure of classified information or any disclosure otherwise prohibited by Federal criminal law, except as otherwise authorized by law.

SEC. 589. ENFORCEMENT, RULEMAKING, AND TECHNICAL STANDARDS.

(a) Federal Trade Commission enforcement.—A violation of section 583, 584, 585, or 586, or of any rule promulgated under any such section, shall be treated as an unfair, deceptive, or abusive act or practice under section 5 of the Federal Trade Commission Act ([15 U.S.C. 45](#)), as amended by this Act. The Commission shall enforce sections 583 through 586, and rules promulgated under such sections, in the same manner, by the same means, and with the same jurisdiction, powers, and duties as though all applicable terms and provisions of the Federal Trade Commission Act were incorporated into and made a part of this subtitle.

(b) State attorney general enforcement.—

(1) In general.—In any case in which the attorney general of a State has reason to believe that an interest of the residents of that State has been or is threatened or adversely affected by a practice that violates section 583, 584, 585, or 586, or a rule promulgated under any such section, the attorney general of the State may bring a civil action on behalf of the residents of the State in an appropriate district court of the United States to—

- (A) enjoin that practice;
- (B) enforce compliance with the applicable section or rule;
- (C) obtain damages, restitution, or other compensation on behalf of residents of the State; or
- (D) obtain such other relief as the court may consider appropriate.

(2) Notice.—Before filing an action under paragraph (1), the attorney general of the State shall provide written notice to the Commission and provide the Commission with a copy of the complaint for such action, except that if it is not feasible to provide such prior notice, the attorney general shall provide such notice immediately upon instituting such action.

(3) Intervention.—Upon receiving notice under paragraph (2), the Commission may intervene in the action, be heard on all matters arising therein, and file petitions for appeal.

(4) Construction.—Nothing in this subsection shall be construed to prevent the attorney general of a State from exercising the powers conferred on the attorney general by the laws of that State to conduct investigations or to administer oaths or affirmations or to compel the attendance of witnesses or the production of documentary and other evidence.

(c) Rulemaking.—

(1) Commission rulemaking.—The Commission shall issue interim final rules to carry out sections 583 through 586 not later than July 1, 2027, and shall thereafter promulgate final rules after notice and comment.

(2) Other rulemaking authority preserved.—Nothing in paragraph (1) shall be construed to limit rulemaking authority otherwise provided under this subtitle to the Secretary of Commerce, the Secretary of Health and Human Services, or any other Federal official.

(d) Technical standards and interoperability.—The Secretary of Commerce, acting through the National Institute of Standards and Technology and in consultation with the Commission and the technical committee established under section 568(e), shall, not later than January 1, 2027, and periodically thereafter, develop, publish, and periodically update technical standards, best practices, conformity profiles, or recognized interoperable standards, as appropriate, to support implementation of this subtitle, including standards or profiles relating to—

- (1) machine-readable provenance data;
- (2) synthetic content labels and other human-readable disclosures;
- (3) preservation of provenance data and labels through editing, export, download, transfer, and republication;
- (4) anti-stripping, anti-falsification, anti-separation, and anti-tampering measures;
- (5) interoperable propagation of provenance data and labels;
- (6) secure methods for associating provenance data and labels with content;
- (7) methods for updating provenance data and labels after material alteration;
- (8) reasonable downstream procedures to reduce removal or defeat of required provenance data, labels, or disclosures;
- (9) channel-specific technical standards for public attention platforms, private communications platforms, app stores, advertising systems, and online marketplaces; and
- (10) such other technical standards, best practices, conformity profiles, or recognized interoperable standards as the Secretary determines necessary to support implementation of this subtitle.

(e) No waiver; no evasion.—

(1) No waiver.—A duty, restriction, prohibition, protection, or remedy established under this subtitle may not be waived, disclaimed, limited, or impaired by contract, agreement, policy, term of service, or other arrangement.

(2) Anti-evasion.—It shall be unlawful for any person to conduct any activity, including by entering into an agreement or contract, engaging in a transaction, structuring an entity, designing or operating a system or interface, or using a technical process, distribution channel, advertising process, ranking system, recommendation system, or other arrangement, to willfully evade or attempt to evade any provision of this subtitle.

(3) Enforcement.—A violation of paragraph (2) shall be enforceable to the same extent and subject to the same civil, administrative, and equitable penalties and remedies as a violation of the provision of this subtitle that the person sought to evade.

(f) Rule of construction.—

(1) Federal floor.—Nothing in this subtitle shall be construed to preempt, displace, or limit any provision of State law that provides equal or greater protection to consumers, minors, vulnerable users, workers, or the public, except to the extent of a direct and irreconcilable conflict with this subtitle.

(2) Other law.—Nothing in this subtitle shall be construed to limit any duty, remedy, penalty, or prohibition otherwise available under Federal or State law.

([The POPULIST Act of 2026](#), TITLE V, version 2.5, last updated August 7, 2026)