

January 5, 2026

RECEIVED

VIA HAND DELIVERY

Jan 05 2026

Anabel Renteria
Initiative Coordinator
Office of the Attorney General
1300 "I" Street, 17th Floor
Sacramento, CA 95814

INITIATIVE COORDINATOR
ATTORNEY GENERAL'S OFFICE

Re: The "California AI Worker Protection Act"

Dear Initiative Coordinator:

On December 1, 2025, I submitted a proposed statewide initiative titled the "California AI Worker Protection Act" ("Initiative") and submitted a request that the Attorney General prepare a circulating title and summary pursuant to section 10(d) of Article II of the California Constitution.

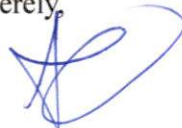
Pursuant to Elections Code section 9002(b), I hereby submit timely amendments to the text of the Initiative. As the proponent of the Initiative, I approve the submission of the amended text to the Initiative and declare that the amendments are reasonably germane to the theme, purpose, and subject of the Initiative. I respectfully request that the Attorney General prepare a circulating title and summary using the amended Initiative.

Please direct all correspondence and inquiries regarding this measure to:

Alexander Oldham

1201 Brickyard Cove Way #218
Point Richmond, CA 94801
510 685-7847
Email: californiaaiact@proton.me

Sincerely,



Alexander Oldham

Enclosures

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

SECTION 1. TITLE, FINDINGS, AND PURPOSE

(a) Title. This measure shall be known as the "California AI Accountability And Worker Protection Act."

(b) Findings. The People of California find:

(1) Frontier artificial intelligence (AI) systems are rapidly increasing in capability and will soon perform complex cognitive work across professional domains, causing significant workforce displacement and other harms to California residents.

(2) These harms are externalities—costs imposed on workers and the public while benefits accrue primarily to AI developers. Workers lose jobs and income. Taxpayers fund social services. Communities experience disruption. Meanwhile, Frontier AI Companies capture the economic benefits.

(3) Basic principles of justice require that those who cause harm bear the cost of that harm, rather than externalizing those costs to innocent parties. This is the principle underlying environmental law, tort law, and workers' compensation.

(4) California has authority under its police power to require businesses conducting research and development in California to prevent, mitigate, or compensate for harms they cause to California residents.

(c) Purpose. This Act establishes a unified framework requiring Frontier AI Companies to identify, prevent, and mitigate the externalized harms caused by their AI systems. Just as environmental law requires those who cause pollution to bear the cost of that pollution, this Act requires those who cause AI-related harms to bear the costs of those harms rather than externalizing them to workers, communities, and the public.

The five categories of externalized harm addressed by this Act—workforce displacement, safety risks, loss of human control, democratic integrity threats, and dangerous concentration of power—represent distinct manifestations of a single phenomenon: the development and deployment of powerful AI systems that create private benefits while imposing costs on others. This Act applies the longstanding legal principle that those who cause harm must bear the cost of that harm across each category of externalized cost.

This Act:

(1) Requires Frontier AI Companies conducting research and development in California to identify and mitigate harms caused by their AI systems;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

- (2) Maintains human control over frontier AI systems;
- (3) Prevents dangerous concentration of power;
- (4) Creates competitive pressure driving all companies toward best practices demonstrated by industry leaders;
- (5) Enables continued innovation by allowing companies with superior practices to advance freely while ensuring laggards do not outpace leaders.

(d) **Constitutional Basis.** This Act exercises California's police power to protect residents from harm. It does not take private property for public use; it requires those who cause harm to bear the cost of that harm. Mitigation obligations are limited by constitutional requirements of nexus and proportionality, which are incorporated into this Act and subject to judicial review.

(e) **Scope Limitation.** This Act regulates frontier AI research, development, and deployment activities occurring in California, and deployment of frontier AI systems to California users. Nothing in this Act shall be construed to regulate activity occurring wholly outside California that has no material connection to California users, California facilities, or California-based personnel.

SECTION 2. DEFINITIONS

(a) "Frontier AI Company" means an entity satisfying all of the following:

(1) *California R&D Requirement.* The entity maintains substantial frontier AI research and development operations in California, including either:

(A) The entity's principal location for frontier AI research and development; or

(B) At least one hundred (100) individuals in California whose primary work involves frontier AI research, development, safety, or deployment.

(2) *Scale Requirement.* The entity meets two or more of the following:

(A) Fair market valuation exceeding ten billion dollars (\$10,000,000,000), where AI constitutes a substantial portion of value or operations;

(B) More than one billion dollars (\$1,000,000,000) in aggregate capital raised for AI development;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(C) More than five hundred million dollars (\$500,000,000) in annual AI research and development expenditure;

(D) More than five hundred million dollars (\$500,000,000) in annual revenue from AI-enabled products or services;

(E) Training of AI systems using compute exceeding 10^{26} floating point operations, or the adjusted threshold set by Commission regulation;

(F) More than five hundred (500) employees worldwide in AI research, development, safety, or deployment.

(3) *Frontier Capability Requirement.* The entity develops or maintains Exclusive Operational Control over AI systems meeting the Frontier Capability Threshold.

(b) "Frontier Capability Threshold" means the level of AI system capability at which systems pose material risk of externalized harms in the categories identified in Section 4(a), as determined by Commission regulation.

(1) *Commission Authority.* The Commission shall define and maintain the Frontier Capability Threshold by regulation, using objective, measurable criteria that identify AI systems posing material risk of:

(A) Workforce displacement at scale affecting California employment;

(B) Safety risks including physical harm, infrastructure compromise, or biosecurity threats;

(C) Loss of effective human control over consequential decisions;

(D) Threats to democratic integrity through manipulation or coordinated inauthentic behavior; or

(E) Dangerous concentration of economic, informational, or political power.

(2) *Required Factors.* In defining the Frontier Capability Threshold, the Commission shall consider and make written findings regarding:

(A) The capability level at which AI systems can substitute for human cognitive labor across professional domains, based on performance on professional licensure examinations, standardized professional tasks, or equivalent measures;

(B) The capability level at which AI systems can operate autonomously over extended periods, adapt to novel circumstances, and execute multi-step plans without human intervention;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(C) Training compute, inference compute, and other technical measures correlated with capability;

(D) Demonstrated real-world effects of AI systems at various capability levels on employment, safety, autonomy, democracy, and concentration;

(E) Projected effects based on capability trajectories and deployment patterns;

(F) Input from the Technical Advisory Committee, affected industries, labor organizations, and the public.

(3) *Form of Definition.* The Commission may define the Frontier Capability Threshold using any combination of:

(A) Performance thresholds on designated evaluation suites or professional examinations;

(B) Technical thresholds such as training compute, parameter count, or inference capabilities;

(C) Functional capability descriptions with objective indicators;

(D) Autonomous operation thresholds;

(E) Other objective, measurable criteria tied to the factors in subdivision (2).

The Commission shall publish specific evaluation methodologies, designated benchmarks, and threshold values. Evaluations used for determining threshold status shall be published and publicly documented before application to any specific company.

(4) *Mandatory Updates.* The Commission shall review and update the Frontier Capability Threshold:

(A) At least annually, with updates taking effect no later than sixty (60) days after publication;

(B) Within ninety (90) days whenever AI systems demonstrate capabilities materially beyond those contemplated in the current definition;

(C) Within ninety (90) days whenever evidence indicates the current threshold fails to capture systems posing material risk of harms in Section 4(a); and

(D) Upon emergency convening under subdivision (5).

(5) *Emergency Update Procedure.* The Commission may convene an emergency review of the Frontier Capability Threshold upon:

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(A) A finding by at least four (4) commissioners that technological developments require immediate threshold revision;

(B) A request by the Governor based on public safety concerns;

(C) Demonstrated capability by any AI system that poses clear and immediate risk of Catastrophic Harm under the current threshold.

Emergency threshold revisions shall take effect upon publication but shall be subject to expedited notice-and-comment within thirty (30) days. Emergency revisions expire after one hundred eighty (180) days unless confirmed through standard rulemaking.

(6) *Interim Threshold.* Until the Commission promulgates its initial definition, and for no longer than one hundred fifty (150) days after the Commission becomes operational, the Frontier Capability Threshold shall be met by any system satisfying one or more of:

(A) Training compute exceeding 10^{26} floating point operations;

(B) Achieving passing scores sufficient for California licensure on professional examinations in three (3) or more of the following: law (California Bar Examination), medicine (USMLE Step 1, 2, and 3), accounting (CPA Examination), engineering (PE Examination), or financial analysis (CFA Examination);

(C) Demonstrating capacity for autonomous operation exceeding forty-eight (48) hours on complex multi-step tasks without human intervention.

The interim threshold shall remain in effect for any company that registered under it until sixty (60) days after the Commission publishes its initial definition.

(7) *Judicial Review.* Any person subject to this Act may challenge the Commission's Frontier Capability Threshold definition as:

(A) Arbitrary, capricious, or unsupported by substantial evidence;

(B) Inconsistent with the factors required by subdivision (2);

(C) Failing to provide fair notice of what systems are covered.

Courts shall review threshold definitions for abuse of discretion, except that challenges based on fair notice shall be reviewed de novo.

(8) *Prospective Application.* Revisions to the Frontier Capability Threshold shall apply prospectively. A company that was not a Frontier AI Company under the prior threshold shall have ninety (90) days from the effective date of a revised threshold to come into compliance if the revision causes the company to meet the definition.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(c) "Covered AI System" means an AI system meeting the Frontier Capability Threshold over which an entity maintains Exclusive Operational Control. An AI system is not a Covered AI System if its model weights have been publicly released under an open license permitting unrestricted use, modification, and redistribution; however, the decision to release weights is itself a Capability Expansion requiring assessment before release.

(d) "Exclusive Operational Control" means control over model weights, training, modification, deployment configuration, and the ability to shut down the system.

(e) "Capability Expansion" means an increase in AI system capabilities, deployment scope, or operational autonomy that materially increases the risk of externalized harms in the categories identified in Section 4(a), as determined by Commission regulation.

(1) *Commission Authority.* The Commission shall define and maintain criteria for Capability Expansion by regulation, identifying the types and magnitudes of changes that constitute Capability Expansion requiring notice or approval under Section 5.

(2) *Required Factors.* In defining Capability Expansion, the Commission shall consider and make written findings regarding:

(A) The relationship between specific capability increases and projected harms in each Harm Category;

(B) The magnitude of capability change that materially increases harm risk, based on empirical evidence and expert assessment;

(C) Deployment scope changes that materially increase exposure of California residents to AI system effects;

(D) Autonomy increases that materially reduce human control or oversight;

(E) Changes to safety measures or guardrails that materially increase harm risk;

(F) The cumulative effect of incremental changes that individually might not constitute Capability Expansion.

(3) *Required Categories.* The Commission's definition shall address, at minimum:

(A) *Compute Increases:* Training or deployment using compute materially exceeding prior systems, with specific thresholds tied to projected capability gains;

(B) *Benchmark Performance Gains:* Performance improvements on designated evaluations exceeding thresholds tied to material capability differences;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(C) *Novel Capabilities*: Demonstrated capabilities not previously exhibited by the company's systems, where such capabilities are on a Commission-maintained Significant Capability List tied to harm categories;

(D) *Deployment Scale*: Increases in California user base or deployment scope exceeding thresholds tied to material increase in aggregate harm exposure;

(E) *Domain Expansion*: Deployment in high-risk domains not previously served, where high-risk domains are designated based on harm potential;

(F) *Autonomy Expansion*: Increases in autonomous operation scope, duration, or consequentiality exceeding thresholds tied to loss-of-control risks;

(G) *Self-Modification*: Any enabling of self-modification, self-replication, or recursive self-improvement;

(H) *Safety Reduction*: Material reduction of safety measures, human oversight, or operational constraints for deployed systems.

(4) *Graduated Treatment*. The Commission may establish graduated categories of Capability Expansion with different procedural requirements based on harm risk:

(A) *Standard Expansion*: Requiring notice and subject to delay under Section 5;

(B) *Lower-Risk Expansion*: Requiring notice with shortened review periods;

(C) *De Minimis Changes*: Exempt from notice requirements under Section 5(l);

(D) *High-Risk Expansion*: Requiring enhanced review, additional assessment, or Commission approval before proceeding.

(5) *Mandatory Updates*. The Commission shall review and update Capability Expansion criteria:

(A) At least annually;

(B) Within ninety (90) days whenever AI capability trajectories indicate current criteria fail to capture material capability increases;

(C) Within ninety (90) days whenever evidence indicates capability increases below current thresholds are causing material harms;

(D) Concurrently with any update to the Frontier Capability Threshold that affects the relationship between capabilities and harms.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(6) *Emergency Updates.* Capability Expansion criteria may be updated through the emergency procedure in subdivision (b)(5), applied mutatis mutandis.

(7) *Interim Criteria.* Until the Commission promulgates its initial definition, Capability Expansion means any of the following:

(A) Training or deploying a system using compute exceeding 10^{26} floating point operations, or fifty percent (50%) greater than the company's previously most capable system, whichever is lower;

(B) Demonstrating performance improvements of twenty percent (20%) or greater on any of the following benchmarks: MMLU, HumanEval, MATH, GSM8K, or successor benchmarks designated by Commission notice;

(C) Deploying to more than five million (5,000,000) new California users within any six (6) month period;

(D) Deploying for the first time in healthcare diagnosis, legal advice, financial advice, critical infrastructure control, weapons systems, or law enforcement applications in California;

(E) Enabling autonomous operation exceeding twenty-four (24) hours without human check-in or approval;

(F) Enabling any self-modification, self-replication, or recursive self-improvement capability;

(G) Removing or materially weakening safety measures, content filters, or human oversight mechanisms present in previously deployed versions.

(8) *Judicial Review.* Capability Expansion criteria are subject to judicial review under the same standards as the Frontier Capability Threshold in subdivision (b)(7).

(f) "California-Derived Revenue" means gross revenue attributable to California users or customers, determined under Commission regulation using reasonable allocation methodologies such as user location, billing address, place of performance, or other objective, verifiable indicia.

(g) "Covered Executive" means any CEO, president, CTO, chief scientist, or equivalent; any director or trustee; any holder of more than five percent (5%) ownership; or any person with practical authority over Covered AI System development, safety, or deployment.

(h) "Commission" means the California AI Safety Commission established by Section 6.

(i) "Harm Categories" means the categories established in Section 4(a).

(j) "Protection Plan" means the plan required by Section 4.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(k) "Catastrophic Harm" means harm of exceptional severity to California, as defined by Commission regulation based on expert assessment of AI risk scenarios.

(1) *Commission Authority.* The Commission shall define Catastrophic Harm by regulation, establishing objective indicators for harms warranting the extraordinary measure of Emergency Orders under Section 5(i).

(2) *Required Categories.* The Commission's definition shall address, at minimum:

(A) *Large-Scale Loss of Life:* Specifying numerical thresholds, categories of mass casualty events, and causal relationships between AI systems and loss of life that constitute Catastrophic Harm;

(B) *Severe Economic Disruption:* Specifying quantitative indicators such as dollar thresholds, percentage of California GDP, unemployment impacts, or market disruption metrics that constitute Catastrophic Harm, distinguishing between acute disruption and gradual economic transition;

(C) *Critical Infrastructure Compromise:* Identifying specific infrastructure categories (power grid, water systems, communications, transportation, financial systems, healthcare systems) and levels of compromise (duration, geographic scope, severity of service disruption) that constitute Catastrophic Harm;

(D) *Democratic Institution Compromise:* Identifying specific institutions (elections, legislative processes, judicial proceedings, regulatory agencies) and types of compromise (manipulation, incapacitation, corruption of information) that constitute Catastrophic Harm, consistent with First Amendment protections.

(3) *Required Factors.* In defining Catastrophic Harm, the Commission shall consider and make written findings regarding:

(A) Expert assessment of plausible AI risk scenarios, including both near-term and longer-term risks;

(B) The severity threshold appropriate for triggering Emergency Orders that override normal procedural protections;

(C) The distinction between catastrophic harms warranting emergency intervention and serious harms addressable through standard enforcement;

(D) The relationship between AI system capabilities and the plausibility of various catastrophic scenarios;

(E) Input from the Technical Advisory Committee, national security experts, infrastructure protection experts, election security officials, and the public.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(4) *Objective Indicators.* For each category of Catastrophic Harm, the Commission shall establish objective indicators that provide clear guidance on when the Emergency Order threshold is met. Indicators shall be:

- (A) Specific enough to provide fair notice and constrain discretion;
- (B) Tied to observable conditions or measurable risks;
- (C) Proportionate to the extraordinary nature of Emergency Order authority.

(5) *Mandatory Updates.* The Commission shall review and update the Catastrophic Harm definition:

- (A) At least annually;
- (B) Within ninety (90) days whenever expert assessment identifies plausible catastrophic risk scenarios not adequately addressed by the current definition;
- (C) Within ninety (90) days whenever a Catastrophic Harm or near-miss occurs that reveals inadequacies in the current definition;
- (D) Concurrently with updates to the Frontier Capability Threshold or Capability Expansion criteria that affect catastrophic risk assessment.

(6) *Emergency Updates.* The Catastrophic Harm definition may be updated through the emergency procedure in subdivision (b)(5), applied mutatis mutandis.

(7) *Interim Definition.* Until the Commission promulgates its initial definition, and for purposes of Emergency Orders under Section 5(i), Catastrophic Harm is limited to circumstances where the Commission has substantial evidence that AI system activity:

- (A) Has directly caused, or is imminently likely to directly cause, loss of ten (10) or more human lives;
- (B) Has directly caused, or is imminently likely to directly cause, physical injury requiring hospitalization to one hundred (100) or more persons;
- (C) Has directly caused, or is imminently likely to directly cause, damage to critical infrastructure as defined by California Government Code Section 8571 that would disrupt essential services to more than one hundred thousand (100,000) California residents for more than twenty-four (24) hours;
- (D) Has directly caused, or is imminently likely to directly cause, compromise of California election systems affecting vote counting, voter registration, or election certification;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(E) Demonstrates that an AI system is operating outside human control in a manner that poses imminent risk of harms in subdivisions (A) through (D).

The interim definition imposes a higher threshold than the Commission's permanent definition may establish, reflecting the need for caution before the Commission has developed full expertise and procedures.

(8) *Judicial Review.* The Catastrophic Harm definition is subject to judicial review under the same standards as the Frontier Capability Threshold in subdivision (b)(7). In addition, any Emergency Order issued under Section 5(i) may be challenged as:

(A) Not supported by substantial evidence of Catastrophic Harm as defined;

(B) Broader than necessary to address the identified catastrophic risk;

(C) Inconsistent with constitutional protections.

Emergency Order challenges shall receive expedited judicial review with initial hearing within seventy-two (72) hours of filing upon request.

(l) "Independent Evaluator" means an entity certified by the Commission under Section 6(k) to conduct assessments, audits, and compliance monitoring, meeting independence, expertise, and confidentiality requirements.

SECTION 3. REGISTRATION

(a) *Registration Required.* Every Frontier AI Company shall register with the Commission within ninety (90) days of this Act taking effect or within sixty (60) days of becoming a Frontier AI Company, whichever is later.

(b) *Contents.* Registration shall include: corporate structure and Covered Executives; description of Covered AI Systems; California operations; current safety and governance practices; and initial Protection Plan.

(c) *Continuation of Operations.* Registration does not require cessation of existing operations. Certification and expansion requirements apply to Capability Expansions occurring after registration.

(d) *Annual Determination.* The Commission shall annually publish a list of Frontier AI Companies and provide a process for entities to request determination of status. Determinations are subject to judicial review.

SECTION 4. HARM CATEGORIES, PROTECTION PLANS, AND MITIGATION

(a) Harm Categories. Frontier AI Companies shall assess and mitigate harms in the following categories. For each category, the harm represents an externalized cost—a cost imposed on workers, the public, or California residents by AI deployment while benefits accrue to the Frontier AI Company. Requiring mitigation of these externalized costs is an exercise of police power, not extraction of benefits.

(1) Workforce Displacement.

(A) The Harm: Economic harm to workers displaced by AI, including lost income, transition costs, unemployment, loss of benefits, and community disruption.

(B) The Externality: Absent mitigation, these costs are borne by displaced workers and California taxpayers while the Frontier AI Company captures automation benefits.

(C) The Obligation: Prevent, mitigate, or compensate for displacement harms caused by the company's AI systems.

(2) Safety Risks.

(A) The Harm: Physical harm to persons, critical infrastructure compromise, biosecurity risks, cybersecurity breaches, and related safety harms.

(B) The Externality: Absent mitigation, safety harms are borne by victims and the public while the company captures deployment benefits.

(C) The Obligation: Implement measures to prevent safety harms and remediate harms that occur.

(3) Loss of Human Control.

(A) The Harm: AI systems operating outside authorized parameters, resisting human oversight, autonomously acquiring capabilities or resources, or evading human control.

(B) The Externality: Loss of control creates risks borne by society while the company may benefit from increased capability.

(C) The Obligation: Maintain effective human control proportionate to system capability and autonomy.

(4) Democratic Integrity.

(A) The Harm: Manipulation of elections, public discourse, or governmental decision-making through AI-generated content, targeted persuasion, or coordinated inauthentic behavior.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(B) **The Externality:** Democratic harms affect all residents but may benefit those seeking to manipulate public opinion.

(C) **The Obligation:** Prevent systems from undermining democratic processes and implement safeguards against misuse.

(5) *Concentration of Power.*

(A) **The Harm:** Excessive concentration of economic, informational, or political power undermining competition, accountability, and distributed decision-making.

(B) **The Externality:** Concentration harms the public interest while benefiting those accumulating power.

(C) **The Obligation:** Maintain distributed governance and prevent deployments creating dangerous concentration.

(b) **Protection Plan Required.** Every Frontier AI Company shall maintain a Protection Plan demonstrating how it prevents, mitigates, or compensates for harms in each category. Plans shall be public except for trade secrets and shall be updated annually and before Capability Expansion.

(c) **Mitigation Principles.**

(1) ***Nexus.*** Mitigation obligations shall have an essential nexus to harms caused by the company's AI systems. Obligations must address identified harms, not unrelated purposes.

(2) ***Proportionality.*** Mitigation obligations shall be roughly proportional to the magnitude and probability of harms caused.

(3) ***No Extraction Without Harm.*** The Commission shall not impose obligations except to address harms caused by the company's AI systems. Obligations not connected to identified harms are prohibited.

(4) ***Causation.*** Companies are responsible for harms caused by their systems. Companies may demonstrate harms are attributable to other factors.

(5) ***Written Findings.*** The Commission shall make written findings establishing the nexus and proportionality of any mitigation obligations or conditions, which findings shall be subject to judicial review.

(6) ***De Novo Proportionality Review.*** Any Frontier AI Company may challenge mitigation obligations as disproportionate to identified harms. Courts shall review proportionality determinations de novo.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(7) *Adjustment for Actual Outcomes.* Mitigation obligations based on projected harms shall be adjusted based on actual outcomes. If projected harms do not materialize, obligations shall be reduced proportionately. If actual harms exceed projections, additional mitigation may be required, subject to nexus and proportionality requirements.

(8) *Workforce Displacement Causation.* Before requiring workforce displacement mitigation, the Commission shall find, based on substantial evidence:

(A) Identifiable displacement in California has occurred or is reasonably certain to occur;

(B) The displacement is attributable to adoption of the company's specific AI systems or systems substantially similar in capability;

(C) The causal contribution of the company's systems to the displacement is not remote or speculative; and

(D) The required mitigation does not exceed the company's proportionate causal contribution to the displacement.

(9) *Limit on Speculative Mitigation.* The Commission shall not require mitigation for projected future workforce displacement unless:

(A) The projection is based on substantial evidence of reasonably foreseeable deployment patterns and workforce impacts;

(B) Mitigation obligations are adjusted retroactively based on actual outcomes; and

(C) If projected displacement does not materialize within three (3) years, prepaid mitigation amounts shall be refunded or credited.

(d) *Flexible Methods.* Frontier AI Companies may satisfy mitigation obligations through any effective method; including:

(1) Prevention through system design or deployment restrictions;

(2) Direct mitigation providing services, support, or resources offsetting harm;

(3) Financial compensation to affected persons or supporting funds;

(4) Direct provision of AI services, goods, or opportunities that offset harms caused;

(5) Other methods demonstrably achieving equivalent mitigation.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

As AI capability increases, direct provision of services by AI systems may satisfy obligations otherwise requiring financial expenditure, if functionally equivalent in mitigating identified harms.

(e) Verification. Effectiveness must be verifiable through measurable outcomes. The Commission may require monitoring, reporting, or third-party verification.

(f) Safe Harbors and Innovation.

(1) The Commission shall establish safe harbors—specific methods that, if implemented, constitute deemed compliance for each Harm Category. Safe harbors provide certainty but are not mandatory.

(2) A Frontier AI Company may propose alternative methods not established as safe harbors. The Commission shall evaluate such methods based on whether they achieve the protected interest, whether effectiveness is verifiable, and how they compare to established methods.

(3) When novel methods prove effective, the Commission may add them to safe harbors.

(g) Evolution. Methods of mitigation shall evolve as AI capabilities evolve. The Commission shall recognize novel methods achieving protected interests through means not previously contemplated and shall update standards as conditions change.

SECTION 5. CERTIFICATION AND CAPABILITY EXPANSION

(a) Certification.

(1) The Commission shall evaluate each Frontier AI Company's Protection Plan and certify companies meeting Demonstrated Best Practices in each Harm Category.

(2) "Demonstrated Best Practices" means practices successfully implemented by at least one Frontier AI Company that represent feasible measures for companies of comparable scale. The Commission shall publish Demonstrated Best Practices at least annually.

(3) Certification is valid for one (1) year and renewable upon continued compliance.

(b) Expansion by Certified Companies. Certified Frontier AI Companies may engage in Capability Expansion upon filing a short-form notice at least fourteen (14) days before expansion. Expansion proceeds unless the Commission issues an emergency order under subdivision (i).

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(c) **Expansion by Non-Certified Companies.** Non-Certified Frontier AI Companies shall file an Expansion Notice at least forty-five (45) days before Capability Expansion, including description of expansion, self-assessment of potential harms, and mitigation commitments. If the Commission does not object within forty-five (45) days, expansion may proceed.

(d) **Harm Assessment for Non-Certified Companies.**

(1) Upon receiving an Expansion Notice, the Commission may conduct a Harm Assessment identifying reasonably foreseeable harms.

(2) The Commission shall provide any draft Harm Assessment to the company, which shall have fifteen (15) days to respond with additional information or objections.

(3) The Commission shall consider the response before finalizing the assessment and determining whether to object.

(4) Any objection or Delay Order shall be based on the finalized assessment and identify specific deficiencies and remediation steps.

(e) **Comparative Delay Orders.** The Commission may issue an Expansion Delay Order to a Non-Certified company only upon written findings of all the following:

(1) The company's Protection Plan is materially inadequate in at least one Harm Category;

(2) At least one Certified company has demonstrated superior practices are feasible for companies of comparable scale;

(3) The deficiency creates material risk of harm; and

(4) Delay is necessary to prevent the company from advancing ahead of Certified competitors with superior protections.

(f) **Delay Order Requirements.** An Expansion Delay Order shall:

(1) Identify specific deficiencies;

(2) Identify which Certified company's practices demonstrate feasibility;

(3) Specify remediation steps;

(4) Set maximum duration not exceeding one hundred eighty (180) days.

The company may terminate delay by demonstrating remediation; the Commission shall act on such demonstration within thirty (30) days. Delay Orders are subject to expedited judicial review.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(g) Commission Shall Not Delay Entire Industry. The Commission may not issue Delay Orders that would delay Capability Expansion by all Frontier AI Companies. If no company has demonstrated superior practices in an area, no company may be delayed for deficiency in that area.

(h) Commission Review Criteria. In evaluating Protection Plans and Expansion Notices, the Commission shall consider:

- (1) The magnitude and probability of each identified harm;
- (2) The effectiveness and credibility of proposed mitigation;
- (3) Demonstrated feasibility based on practices implemented by other Frontier AI Companies;
- (4) Proportionality between expansion benefits and residual harms;
- (5) The company's track record of compliance and mitigation performance.

(i) Emergency Orders. Upon substantial evidence of clear and present danger of Catastrophic Harm, imminent loss of human control, or active circumvention of required safety controls, the Commission may issue an emergency order taking immediate effect. Emergency orders expire within fourteen (14) days unless the Commission conducts an evidentiary hearing and issues written findings to extend for up to sixty (60) days. Affected parties may seek expedited judicial review.

(j) Timeline Protections.

(1) If the Commission fails to act on an Expansion Notice within specified timelines, the Frontier AI Company may proceed with the noticed expansion or seek expedited judicial review.

(2) For Capability Expansions designated as lower-risk by Commission regulation, failure to object within thirty (30) days constitutes approval.

(k) Certification Modification and Revocation.

(1) *Modification.* The Commission may modify Certification conditions upon finding material change in circumstances. Before modification, the Commission shall provide written notice of proposed modifications and thirty (30) days to respond. Modifications shall comply with nexus and proportionality requirements.

(2) *Revocation.* The Commission may revoke Certification upon finding material noncompliance or violation. Except in emergencies under subdivision (i), revocation requires written notice of grounds and opportunity for hearing. Revocation is subject to expedited judicial review.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(l) *De Minimis Exception.* Actions otherwise constituting Capability Expansion are exempt if the company reasonably determines and documents in writing that the action does not materially increase risk in any Harm Category. Such determinations are subject to Commission audit; erroneous determinations made in bad faith constitute violations.

(m) Firm Timeline Protections and Remedies.

(1) *Deemed Approval.* If the Commission fails to issue an objection, Delay Order, or request for additional information within the timelines specified in this Section, expansion is deemed approved.

(2) *Expedited Court Relief.* If the Commission exceeds statutory timelines, the Frontier AI Company may petition the Superior Court for the County of Sacramento for an order deeming the expansion approved. The court shall grant the petition upon finding the Commission exceeded timelines absent good cause.

(3) *No Constructive Denial.* The Commission shall not effectively deny or delay expansions through requests for excessive information, unreasonable conditions, or procedural obstacles. A Frontier AI Company may petition for relief under this subdivision if Commission conduct effectively prevents timely expansion despite nominal compliance with timelines.

SECTION 6. CALIFORNIA AI SAFETY COMMISSION

(a) *Establishment.* There is established the California AI Safety Commission as an independent body in the executive branch.

(b) *Powers.* The Commission may: register Frontier AI Companies; evaluate Protection Plans and issue Certifications; approve, condition, or delay Capability Expansions; certify Independent Evaluators; investigate violations; impose civil penalties; adopt regulations; subpoena evidence; conduct audits of California operations; coordinate with federal and international authorities; and take actions necessary to implement this Act. Commission orders and requirements shall be narrowly tailored to the Act's purposes and to activities within California's regulatory jurisdiction.

(c) *Composition.* Seven (7) commissioners:

(1) Three (3) appointed by the Governor, subject to Senate confirmation;

(2) Two (2) appointed by the Attorney General;

(3) One (1) appointed by the Senate Committee on Rules;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(4) One (1) appointed by the Speaker of the Assembly.

(d) Qualifications.

(1) Each commissioner shall have demonstrated expertise in one or more of: AI technology; AI safety and alignment; economics of technological change and workforce transitions; long-term implications of transformative technologies; corporate governance; or technology regulation.

(2) Each commissioner shall have a demonstrated track record of:

(A) Considering long-term and systemic consequences of technology;

(B) Orientation toward the public interest; and

(C) Serious engagement with transformative AI issues.

(3) Before assuming office, each commissioner shall execute a public written statement affirming understanding of the Act's purposes and commitment to prioritizing long-term public interests over regulated entity interests.

(e) Independence Requirements.

(1) *Disqualified.* No person may serve who:

(A) Is currently employed by or consulting for any Frontier AI Company;

(B) Holds financial interests exceeding two hundred fifty thousand dollars (\$250,000) in any single Frontier AI Company, excluding diversified funds;

(C) Has an immediate family member who is a Covered Executive of any Frontier AI Company; or

(D) Within the preceding three (3) years, was employed by any Frontier AI Company in a senior role or received compensation exceeding one hundred thousand dollars (\$100,000) from Frontier AI Companies in any year.

(2) *May Serve with Disclosure and Recusal.* A person may serve if:

(A) Employed by a Frontier AI Company more than three (3) years before appointment, with disclosure, divestiture of interests exceeding fifty thousand dollars (\$50,000), and recusal from that company's matters for two (2) years;

(B) Employed by or advising AI companies that are not Frontier AI Companies, with disclosure;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(C) Holding interests under fifty thousand dollars (\$50,000) in Frontier AI Companies, with disclosure and recusal on that company's matters; or

(D) Holding diversified funds including AI stocks.

(3) *Disclosure.* Commissioners shall file comprehensive disclosure of AI-related employment, compensation, and financial interests within ten (10) years, updated annually and made public.

(f) *Terms and Removal.* Commissioners serve staggered six (6) year terms, may be reappointed once, and may be removed only for cause (neglect, malfeasance, bias favoring regulated entities, or incapacity) after written notice and opportunity to respond.

(g) *Post-Service Restrictions.* For five (5) years after service, former commissioners shall not: accept employment with any Frontier AI Company; accept compensation exceeding ten thousand dollars (\$10,000) annually from any Frontier AI Company; or represent any Frontier AI Company before the Commission. Violation is a misdemeanor and shall result in disgorgement of compensation.

(h) *Public Confirmation.* Governor appointments subject to Senate confirmation shall include public hearings examining qualifications and commitment to the Commission's mission.

(i) *Executive Director and Staff.* The Commission shall appoint an Executive Director meeting independence requirements and employ staff with expertise in AI technology, AI safety, labor economics, enforcement, and law. Staff salaries may exceed standard state levels for specialized AI expertise to attract qualified personnel. Staff are subject to independence requirements and post-service restrictions established by Commission regulation.

(j) *Advisory Committees.* The Commission shall establish standing advisory committees—Technical, Economic and Workforce, and Governance and Ethics—composed of experts providing non-binding advice.

(k) *Independent Evaluators.* The Commission shall certify Independent Evaluators meeting independence, expertise, and confidentiality requirements to assist with assessments and audits under Commission supervision. Certification requires demonstrated expertise, independence from Frontier AI Companies, adequate confidentiality protections, and absence of conflicts.

(l) *Confidentiality and Trade Secret Protection.*

(1) The Commission shall implement strong protections for confidential business information and trade secrets, including: protective orders limiting access and use; secure facilities and systems; access limited to personnel with need to know; and prohibition on disclosure except as required for Commission functions.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(2) Unauthorized disclosure by Commission personnel or contractors is a misdemeanor punishable by up to one (1) year imprisonment and shall result in termination and civil liability.

(3) Trade secret information shall not be disclosed publicly except in summary or redacted form not revealing proprietary details.

(m) California AI Safety Fund. There is established the California AI Safety Fund for exclusive use in administering this Act.

(1) *Sources*: Initial Registration Fees; annual assessments; civil penalties; other fees; interest; and gifts.

(2) *No General Fund Impact*. The Commission shall be funded entirely from regulated entities.

(n) Initial Registration Fees. Upon registration:

(1) Companies with California-derived revenue exceeding one billion dollars: five million dollars (\$5,000,000);

(2) Revenue between one hundred million and one billion dollars: one million dollars (\$1,000,000);

(3) Revenue below one hundred million dollars: two hundred fifty thousand dollars (\$250,000).

Initial fees are credited against first annual assessment. Until the Commission is operational, the Attorney General shall collect fees and deposit them in the Fund.

(o) Annual Assessments. The Commission shall collect annual assessments:

(1) For companies with California-derived revenue: a percentage not exceeding one-half of one percent (0.5%) of such revenue;

(2) For pre-revenue companies: assessment based on California AI R&D expenditures under Commission methodology;

(3) Minimum assessment: fifty thousand dollars (\$50,000);

(4) Maximum assessment: twenty million dollars (\$20,000,000).

(p) Temporary Advance. If initial fees are insufficient to establish the Commission, the Controller may advance up to ten million dollars (\$10,000,000) from the General Fund, to be repaid with interest within three (3) years from first annual assessments.

SECTION 7. ENFORCEMENT

(a) Civil Penalties. The Commission may impose civil penalties considering proportionality, culpability, harm, duration, and cooperation:

(1) Capability Expansion in violation of Delay or Emergency Order: up to twenty percent (20%) of California-derived revenue during violation period;

(2) Material violation of Protection Plan or Certification conditions: up to fifteen percent (15%) of California-derived revenue during violation period;

(3) Material misrepresentation to Commission: up to ten percent (10%) of California-derived revenue;

(4) Failure to file required reports: up to fifty thousand dollars (\$50,000) per day;

(5) Other violations: up to fifty million dollars (\$50,000,000), proportionate to severity.

(b) Executive Accountability. Covered Executives who knowingly authorize or fail to prevent violations are personally liable for up to one million dollars (\$1,000,000) per violation and disgorgement of compensation during violation period.

(c) Criminal Penalties. The following, when done willfully and knowingly, are felonies punishable by two (2) to six (6) years imprisonment:

(1) Capability Expansion violating Commission Emergency Order where expansion creates substantial risk of Catastrophic Harm;

(2) Knowingly making materially false statements to Commission regarding capabilities or safety mechanisms;

(3) Knowingly destroying or concealing evidence after Commission preservation demand.

(d) Private Enforcement. California residents, Frontier AI Company employees, or public-interest organizations may bring civil actions for declaratory or injunctive relief for ongoing violations, subject to pre-suit notice and thirty (30) day cure period. Prevailing plaintiffs recover reasonable attorney's fees.

(e) Whistleblower Protections. No retaliation for reporting violations or participating in enforcement. Remedies include reinstatement, back pay, compensatory damages, and attorney's fees. Whistleblowers providing original information leading to penalties exceeding one million dollars (\$1,000,000) receive fifteen to twenty-five percent (15-25%) of penalties collected.

SECTION 8. JURISDICTION AND SCOPE

(a) **Scope.** This Act applies to Frontier AI Companies as defined in Section 2(a). Commission orders and requirements shall be narrowly tailored to regulate:

(1) Frontier AI research, development, training, evaluation, modification, governance, and deployment activities occurring in California; and

(2) Deployment of Covered AI Systems to California users.

The Commission shall not issue orders regulating activity occurring wholly outside California except to the extent such activity directly affects California users or is conducted using California facilities or personnel.

(b) **Anti-Evasion.** Transactions designed primarily to evade this Act while maintaining California operations may be disregarded for jurisdictional purposes. Entities may not avoid coverage by restructuring if combined operations would meet the definition.

(c) **Notice of Material Changes.** Frontier AI Companies shall provide ninety (90) days' notice before relocating principal AI R&D operations from California or transferring material AI assets to entities not subject to this Act. The Commission may impose conditions on such transfers to ensure continued protection of California residents.

(d) **Subsidiaries.** Where a Frontier AI Company is part of a larger entity, all affiliates engaged in frontier AI development are jointly responsible for compliance.

(e) **Registered Agent.** Each Frontier AI Company shall designate an agent for service of process in California.

(f) **California-Specific Implementation.**

(1) Nothing in this Act shall be construed to require a Frontier AI Company to modify the operation, capabilities, training, or architecture of an AI system as deployed to users outside California.

(2) Frontier AI Companies may satisfy requirements of this Act through California-specific measures, including but not limited to:

(A) California-specific deployment configurations, access controls, or user interface modifications that implement required safeguards for California users without affecting service to non-California users;

(B) California-specific monitoring, auditing, and reporting systems;

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(C) California-specific incident response and harm mitigation protocols;

(D) Any other technically feasible means of implementing requirements for California users or operations without affecting non-California service.

(3) The Commission shall not find a Protection Plan or mitigation measure inadequate solely because it is implemented through California-specific measures rather than system-wide changes, provided the California-specific measures achieve equivalent protection for California residents.

(4) Where a Frontier AI Company demonstrates that a required measure cannot be implemented for California users without affecting non-California users, the Commission shall consider technically feasible alternatives that achieve the protected interest while minimizing extraterritorial effect.

(5) *Limitation on Training Restrictions.* The Commission shall not condition approval of Capability Expansion on modifications to the training process of AI systems, except:

(A) Training occurring in California using California infrastructure; or

(B) Training using data obtained from California users in violation of applicable California law.

(g) *No Requirement of Global Uniformity.* A Frontier AI Company shall not be required to implement identical policies, safeguards, capabilities, or restrictions globally as a condition of compliance with this Act. The Commission shall evaluate compliance based on protections afforded to California users and activities occurring in California, not on the company's global policies or deployments.

SECTION 9. GENERAL PROVISIONS

(a) *Severability and Priority.*

(1) *General Severability.* If any provision of this Act or application thereof is held invalid, the remainder continues in full force and effect.

(2) *Independence of Harm Categories.* Each Harm Category in Section 4(a) is independent. If any Harm Category is held invalid or unenforceable, the remaining categories continue in full force and effect.

(3) *Priority of Provisions.* If any provision is held unconstitutional but could be modified to cure the deficiency, courts are requested to modify rather than invalidate. If modification is

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

not possible, courts are requested to invalidate the minimum necessary to cure the constitutional deficiency.

(4) *Core Provisions.* The People identify the following as core provisions that shall survive even if other provisions are invalidated:

- (A) The registration requirement (Section 3);
- (B) The Commission establishment (Section 6(a)-(l));
- (C) The safety risk harm category (Section 4(a)(2));
- (D) The enforcement mechanisms (Section 7).

(5) *Declaration.* The People of California declare they would have adopted this Act and each provision, section, subdivision, sentence, clause, and phrase thereof regardless of whether any other provision was declared invalid.

(b) *Construction.* This Act shall be liberally construed to effectuate its purposes consistent with constitutional limits.

(c) *Rulemaking.* The Commission shall adopt regulations implementing this Act following California APA procedures, including:

- (1) Frontier Capability Threshold criteria and designated evaluations;
- (2) Capability Expansion criteria and graduated treatment categories;
- (3) Catastrophic Harm definition and objective indicators;
- (4) Demonstrated Best Practices standards and safe harbors;
- (5) Harm Assessment methodologies with objective indicators;
- (6) Expansion evaluation procedures and timelines;
- (7) Lower-risk expansion criteria for expedited review;
- (8) Assessment calculations and California-derived revenue methodology;
- (9) Independent Evaluator certification requirements;
- (10) Trade secret protection procedures.

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(d) **Mandatory Review.** The Commission shall comprehensively review and update regulations:

(1) At least annually for all technical definitions including the Frontier Capability Threshold, Capability Expansion criteria, and Catastrophic Harm definition;

(2) At least every three (3) years for all other regulations;

(3) Within one hundred eighty (180) days whenever:

(A) AI systems demonstrate capabilities materially beyond those contemplated in current definitions;

(B) Evidence indicates current thresholds fail to capture systems or changes posing material risk of harms in Section 4(a);

(C) Evidence indicates current thresholds capture systems or changes not posing material risk of harms in Section 4(a);

(D) Traditional employment or economic metrics become materially inapplicable to measuring protected interests;

(E) A majority of Frontier AI Companies propose mitigation methods materially different from current safe harbors;

(F) The Commission determines existing frameworks inadequate to address identified harms; or

(G) A significant AI-related incident in California or elsewhere reveals gaps in current regulatory frameworks.

(4) The Commission shall publish an annual report assessing whether current definitions remain appropriately calibrated to technological capabilities and harm risks, including:

(A) Analysis of capability trajectories and projected near-term developments;

(B) Assessment of whether current thresholds are capturing the appropriate systems;

(C) Review of harm data and whether projected harms are materializing;

(D) Recommendations for threshold adjustments.

(e) **First Amendment Protections.**

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

(1) Nothing in this Act shall be construed to prohibit, restrict, or require approval for the publication, dissemination, or communication of ideas, opinions, information, or expression protected by the First Amendment to the United States Constitution or Article I of the California Constitution.

(2) The capability expansion provisions of this Act regulate the development and deployment of AI systems as commercial products and services, not the content of expression generated by such systems. Required approvals, assessments, and conditions relate to the system's capabilities, safety characteristics, and deployment scope—not to the viewpoint, content, or message of any expression.

(3) The democratic integrity provisions of Section 4(a)(4) address deceptive conduct and inauthentic coordination in connection with elections, not the regulation of protected political speech. The Commission shall interpret and apply these provisions consistent with First Amendment jurisprudence regarding fraud, deception, and election integrity.

(4) *Safe Harbor for Expressive Platforms.* AI systems deployed primarily as platforms for user expression, where the AI system facilitates rather than generates the expression, shall not be subject to capability expansion requirements based solely on the content of user expression facilitated by the platform.

(5) Any Commission order that would operate as a content-based restriction on protected expression shall be subject to strict scrutiny upon judicial review.

(f) Relationship to Federal Law.

(1) *Complementary Regulation.* This Act is intended to complement, not conflict with, federal AI regulation. Where federal law establishes minimum standards for AI development, deployment, or safety, this Act may impose additional requirements not inconsistent with federal standards.

(2) *Federal Floor.* If federal law establishes requirements in areas covered by this Act, the requirements of this Act shall apply to the extent they exceed federal requirements but shall not apply to the extent they would conflict with or frustrate federal requirements.

(3) *Express Preemption.* This Act shall be preempted only by federal law that expressly states an intent to preempt state AI regulation or that creates a direct conflict that makes compliance with both state and federal law impossible.

(4) *Cooperation.* The Commission shall seek to coordinate with federal agencies regulating AI and, where appropriate, enter into cooperative agreements for information sharing, joint enforcement, and regulatory harmonization.

(g) Amendment. This Act may be amended by two-thirds vote of each house of the Legislature to further its purposes. Amendments narrowing coverage, weakening

CALIFORNIA AI ACCOUNTABILITY AND WORKER PROTECTION ACT

enforcement, reducing Commission authority, or limiting mitigation requirements shall not take effect without voter approval.

(h) **Effective Date.** This Act takes effect the day after voter approval. The Commission shall be operational within one hundred eighty (180) days. Until operational, the Attorney General shall receive registrations, collect fees, and may seek emergency relief for imminent Catastrophic Harm from California-based AI activities.