

Artificial Intelligence


Risk-Based Regulatory Framework Permits Safe Innovation of AI

The California Chamber of Commerce takes responsible and safe innovation of AI seriously and generally shares in the Legislature’s and Governor’s overarching goal in promoting reasonable safeguards in AI innovation. Unfortunately, the only 2024 legislation on the matter, SB 1047 (Wiener; D-San Francisco), went far beyond establishing safeguards, seeking to make certain developers guarantee that their models would never result in certain harms, subject to significant liability.

While the bill was a moving target, constantly changing requirements and making it difficult to analyze the full impact of the legislation, in the end, the CalChamber’s major concerns, including concerns around liability, open source, the impact on the AI ecosystem, unworkable requirements such as full shutdowns, cloud compute, and more, remained unaddressed.

Although the CalChamber agrees that regulatory efforts to promote AI safety are critical, [SB 1047 missed the mark](#) entirely in how it chose to get there, fixating on demanding unrealistic guarantees, imposing untenable liability risks regardless of culpability, prescribing extremely intrusive and industry-killing “know your customer” requirements, as well as kill switches and full shutdown mandates. Unfortunately, there was no amount of fixing or fine tuning the bill that would have adequately addressed concerns because the bill was broken at its core in the approach it chose to take.

The CalChamber instead would support reasonable safety frameworks that do not regulate the AI systems or models



themselves. The focus should be on requiring certain best practices and/or prohibiting certain applications and punishing bad actors, not regulating the development of the actual technology.

ISSUE

At the Joint California Summit on Generative AI held at the University of California, Berkeley in May 2024, Governor Gavin Newsom aptly summarized what is at risk with AI regulations when he stated: “if we over-regulate, if we overindulge, if we chase a shiny object, we could put ourselves in a perilous position.”

No bill better embodied that statement than SB 1047 in 2024. Regulating a technology that does not yet exist, for threats that in no way appear to be imminent, over the objections of the widest range of stakeholders to have banded together on any single AI bill to warn about the perils that will befall the AI ecosystem, is confounding at best. From a safety standpoint, from a technological innovation standpoint, and from an economic standpoint, California cannot afford to get this wrong.

At several points, the author of the bill pointed to congressional inaction on any number of issues from social media to data privacy, drawing false equivalencies to this policy issue to justify forcing the policy forward. First, these are global issues warranting federal solutions. Second, the federal government not only has a responsibility to act, but they also are actively

taking action. Fracturing the regulatory landscaping and undermining federal efforts does not make California any safer. In fact, it does the opposite.

Many CalChamber members have actively supported Governor Newsom’s Executive Order and the Biden White House Executive Order, as well as the White House voluntary commitments, and other voluntary commitments around the world to help move toward safe, secure, and transparent development of AI technology, because they support these goals. Along these lines, the CalChamber is open to supporting similar commitments via an executive order or bill in 2025, perhaps building on the safety standards that were just released by the National Institute of Standards and Technology (NIST) U.S. AI Safety Institute pursuant to the White House executive order.

RUNNING RISK OF MAKING CALIFORNIA MORE VULNERABLE TO GLOBAL THREATS, UNDERMINING ECONOMIC AND TECHNOLOGICAL INNOVATION

The importance of ensuring consistency in the AI regulatory landscape nationally, and the need to follow federal guidance on certain issues that transcend national borders cannot be overstated. However well-intentioned, SB 1047 would have done precisely what the business community has warned against doing when legislating AI: regulating the technology itself, threatening California’s footing as the home of the world’s leading AI companies. By weakening California’s competitive advantage, SB 1047 would have opened the door for other countries to dominate the future of AI — countries that may not play by the same rules that SB 1047 sought to force upon developers in California.

Regulatory inconsistency and uncertainty, high compliance costs, and significant liability risks imposed on developers for failing to guarantee against harmful uses of their models by third parties ultimately will have a dramatic and potentially devastating impact on the entire AI ecosystem, discouraging economic and technological innovation. Instead of making Californians safer, the bill would only hamstring businesses from developing the very AI technologies that could protect against dangerous models developed elsewhere.

RISKING DISRUPTION, DEVASTATION OF ENTIRE AI ECOSYSTEM

During an incredibly challenging budget year, SB 1047 risked significant costs to the state in the realm of tens of millions of dollars just in terms of the incredible potential for future tax revenue that the AI ecosystem can bring to California alone

— meaning, not simply from AI companies, but also from all the industries and businesses looking to leverage AI to increase their efficiency and profitability.

Again, it would be a mistake to enact legislation that regulates the development of technology itself instead of the implementation and uses of it. Such legislation creates a hostile environment for innovation and drives investment to other tech hubs, both inside and outside the United States, with far-reaching implications for state revenues.

Even if AI legislation such as SB 1047 seeks to target only “Big Tech,” SB 1047 demonstrated how the realistic impact of that legislation may not be so limited. AI startups, small businesses, researchers, independent labs, academics, and federal policy experts all spoke out against SB 1047, detailing the ways in which their own interests would have been hurt.

These are entities that stand to lose the possibility of building on the latest, more capable AI models in order to enter into the market or to stay competitive in the market. These are entities that rely on access to those models to apply them toward society’s biggest challenges. Interestingly, they also are entities that often do not all align on the same side of an issue.

Even after numerous amendments, in the end, SB 1047 merely touched on certain problems on the periphery of the bill, such as the removal of a penalty of perjury. On the whole, the amendments failed to address the vast majority of the concerns, including that the bill (1) placed untenable liability risks on developers and effectively foreclosed open-sourcing large models; (2) imposed an intrusive and unreasonable Know Your Customer Obligation and kill switch requirements; and (3) created regulatory uncertainty, suffering from vagueness issues as well as overbreadth.

CalChamber Concerns with SB 1047

- First, **CalChamber supports holding bad actors accountable for their bad acts — which existing law already does.** Unfortunately, that was not what SB 1047 did. Instead, it would have held developers liable for any potential harm caused by a model built off their original model, even if they had no role in building that other model and regardless of the acts of intervening third parties. For instance, a third party could fine tune a model on Chemical, Biological, Radiological, and Nuclear (CBRN) data that the original developer did not. Yet the original developer is being asked to make guarantees about what the third party may or may not do, years, if not decades, down the line.

Imagine requiring designers or developers of engines of a certain horsepower to guarantee that no one can use or

misuse the engine to build a car or other product developed in the future that would be unreasonably dangerous, and then holding them automatically liable for any resulting harm from the end product, even if the engine component was not defective and they had no role in developing the end product.

- Second, **the bill imposed significantly problematic obligations on operators of computing clusters** (for example, data centers or companies that provide cloud computing for frontier model training), requiring them to collect personally identifiable data from their prospective customers, predict if a prospective customer “intends to utilize the computing cluster to deploy a covered model,” and then implement a kill switch to enact a full shutdown in an emergency. These obligations violate customer privacy and security, creating significant risk that customers will move away from U.S.-based cloud providers.

- Finally, among the many examples of the regulatory uncertainty and vagueness or overbreadth issues, were the definitions of “critical harm,” “reasonable care” and “covered model.” Specifically, “critical harms” was so broad that it would have included not only weapons of mass destruction, but also automated phishing campaigns. And when mandating “reasonable care” in the context of speculative CBRN risks, it was unclear in what scenario it might ever be reasonable to move forward with a model if a developer could not totally eliminate the possibility of a critical harm based on future intervening acts of a third party.

By using *computing power* and cost, rather than *capability*, to define covered models, the bill equated model size/cost to risk and managed to be simultaneously both overly broad and too narrow. That means critical harms caused by less costly and more efficient AI models can continue to be developed, unchecked.

In the end, such deficiencies were much more likely to hamstring developers from innovating the technologies that can protect Californians and discourage the growth of the AI economy in a state that currently houses 35 of the 50 leading AI companies in the world, 21 of them in San Francisco.

GOVERNOR'S VETO MESSAGE

In vetoing SB 1047, Governor Newsom stated:

“California is home to 32 of the world’s 50 leading AI companies, pioneers in one of the most significant technological advances in modern history. We lead in this space because of our research and education institutions, our diverse and motivated workforce, and our free-spirited cultivation of intellectual freedom. As

stewards and innovators of the future, I take seriously the responsibility to regulate this industry. [...]

“SB 1047 magnified the conversation about threats that could emerge from the deployment of AI. Key to the debate is whether the threshold for regulation should be based on the cost and number of computations needed to develop an AI model, or whether we should evaluate the system’s actual risks regardless of these factors. [...]

“By focusing only on the most expensive and large-scale models, SB 1047 establishes a regulatory framework that could give the public a false sense of security about controlling this fast-moving technology. Smaller, specialized models may emerge as equally or even more dangerous than the models targeted by SB 1047 — at the potential expense of curtailing the very innovation that fuels advancement in favor of the public good.

“Adaptability is critical as we race to regulate a technology still in its infancy. This will require a delicate balance. While well-intentioned, SB 1047 does not take into account whether an AI system is deployed in high-risk environments, involves critical decision-making or the use of sensitive data. Instead, the bill applies stringent standards to even the most basic functions — so long as a large system deploys it. I do not believe this is the best approach to protecting the public from real threats posed by the technology.

“Let me be clear — I agree with the author — we cannot afford to wait for a major catastrophe to occur before taking action to protect the public. California will not abandon its responsibility. Safety protocols must be adopted. Proactive guardrails should be implemented, and severe consequences for bad actors must be clear and enforceable. I do not agree, however, that to keep the public safe, we must settle for a solution that is not informed by an empirical trajectory analysis of AI systems and capabilities. Ultimately, any framework for effectively regulating AI needs to keep pace with the technology itself.

“To those who say there’s no problem here to solve, or that California does not have a role in regulating potential national security implications of this technology, I disagree. A California-only approach may well be warranted — especially absent federal action by Congress — but it must be based on empirical evidence and science. The U.S. AI Safety Institute, under the National Institute of Science and Technology, is developing guidance on national security risks, informed by evidence-based approaches, to guard against demonstrable risks to public safety. Under an Executive Order I issued in September 2023, agencies within my Administration are performing risk analyses of the potential threats and vulnerabilities to California’s critical infrastructure using AI. These are just a few examples of the many endeavors

underway, led by experts, to inform policymakers on AI risk management practices that are rooted in science and fact. And endeavors like these have led to the introduction of over a dozen bills regulating specific, known risks posed by AI, that I have signed in the last 30 days.

“I am committed to working with the Legislature, federal partners, technology experts, ethicists, and academia, to find the appropriate path forward, including legislation and regulation. Given the stakes — protecting against actual threats without unnecessarily thwarting the promise of this technology to advance the public good — we must get this right.”

GOVERNOR’S WORKING GROUP

On the same day he vetoed SB 1047, Governor Newsom issued a [press release](#) announcing a new working group, building on the partnership created after his 2023 executive order. Among the working group members are leading experts on GenAI, including the “godmother of AI,” Dr. Fei-Fei Li, as well as Tino Cuéllar, member of the National Academy of Sciences Committee on Social and Ethical Implications of Computing Research, and Jennifer Tour Chayes, dean of the College of Computing, Data Science, and Society at UC Berkeley. Stating that “[w]e have a responsibility to protect Californians from potentially catastrophic risks of GenAI deployment,” the Governor made clear that they will both quickly and thoughtfully move toward “a solution that is adaptable to this fast-moving technology and harnesses its potential to advance the public good.” To that end, the Governor asked the group to help California develop workable guardrails for deploying GenAI, focusing on developing an empirical, science-based trajectory analysis of frontier models and their capabilities and attendant risks.

On December 11, 2024, the group issued an update, stating that it is beginning its work by developing a draft report that

draws upon academics and experts from a variety of disciplines and is anticipated to be shared in the first quarter of 2025. That report will include a review of recent literature and research, outlining the latest scientific understanding about frontier model capabilities and risks. To encourage feedback and input from a wide range of expertise, they plan to convene a series of stakeholder activities which may include structured workshops, in-person or remote sessions, and an open opportunity for interested parties to submit written comments about the topics covered in the draft report. That feedback will be incorporated, and they will publish a final report for the Governor’s and Legislature’s consideration, anticipated by summer 2025.

At the same time, the working group indicated that they will facilitate an open call for additional comments, reflections, and ideas for partnership moving forward to “further advance scholarship and multi-sector collaboration”. (See [Update from the Co-Leads of the Joint California Policy Working Group on AI Frontier Models](#).)

CALCHAMBER POSITION

Ultimately, as with all AI legislation, any regulatory framework should be risk-based and avoid regulating the technology itself. Accordingly, the CalChamber is open to supporting reasonable safety frameworks that do not regulate the AI systems or models themselves, or supporting voluntary commitments via an executive order or bill in 2025.

Recognizing, however, that the Governor has expressed a clear intention to take a thoughtful — yet swift — approach to issuing workable guardrails that will be informed by his working group experts, and sharing in that same goal of supporting reasonable and workable guardrails, the CalChamber will engage in any processes and stakeholder opportunities made available by that working group.



Staff Contact
Ronak Daylami
Policy Advocate

ronak.daylami@calchamber.com

January 2025