April 14, 2026

TO: Members, Senate Judiciary Committee

SUBJECT: SB 1119 (PADILLA) COMPANION CHATBOTS: CHILDREN'S SAFETY
OPPOSE UNLESS AMENDED – AS AMENDED MARCH 25, 2026
SCHEDULED FOR HEARING – APRIL 21, 2026

The California Chamber of Commerce and the undersigned respectfully must **OPPOSE UNLESS AMENDED SB 1119 (Padilla)** as amended March 25, 2026. The bill is an important step in the right direction and preserves one of the most important topics of discussion within the Legislature where it belongs. The Legislature is where all stakeholders can have a seat at the table and discuss how to strike the right balance on weighty issues, in stark contrast to what we saw over the interim with ballot initiative proposals that would have limited the role of this institution and the participation of key stakeholders. We hope that you will continue to provide the business community a meaningful opportunity to participate in the conversations around shaping what this law should look like. Our intent is to help strike the balance between implementing age-appropriate guardrails for kids to protect our most vulnerable population from foreseeable harms – a responsibility we all have to take seriously and prioritize – and ensuring that we use the best tools available to educate and prepare our youth for a future that will undoubtedly require fluency with a wide range of technologies.

Technological innovation is the backbone of California, not only creating jobs and economic opportunity, but driving the research behind medical and scientific breakthroughs, and ensuring the state remains a global hub for social and economic progress and shared prosperity—including for those children.

We appreciate the general approach that **SB 1119** takes by allowing youth access to AI tools while ensuring that there are different levels of protection applied to their experiences through impact assessments, default protections, and parental controls. We also appreciate that the bill aligns with AB 1043 (Wicks, Ch. 675, Stats. 2025), the Digital Age Assurance Act, instead of mandating that companies create some new technology or mechanism to estimate age. Leveraging an existing framework removes the unnecessary burden and cost for businesses to build a second, duplicative (if not potentially divergent) verification system. (*See* Proposed Section 22611.)

That being said, we also hope that you would consider that the thrust of **SB 1119** overlaps with SB 243 (Padilla, Ch. 677, Stats. 2025) which was just signed into law last fall. That bill is currently being replicated across other states in the nation and only goes into effect starting July 1, 2027, in California. Of course, we agree that these conversations are still important to have, but we urge you to consider taking a longer view of what should be done in this area, so we have the benefit of learning from the implementation of SB 243. Nonetheless, we look forward to engaging in these conversations with you on **SB 1119**, and have identified several broad areas of concern for your consideration as we continue to work through these issues and providing suggestions on potential solutions.

Scope

While **SB 1119** largely appears to be addressing concerns related to "children's" use of certain AI tools, certain provisions within the bill on their face appears to apply to adults by way of reference to "users", which is not defined in the bill in contrast to "child" which is specifically defined to mean anyone under 18 years of age. This suggests obligations applying to "users" are intended to apply to all users, regardless of age, not just minors. A bill titled "companion chatbots: children's safety" should not introduce new obligations for adult users without it being much clearer in order to give impacted businesses fair warning of what the law requires. If the bill intends to apply to adults, it should be restructure and retitled to accordingly to make it clear which obligations apply to which age groups.

Defining "harm"

Perhaps one of the most important, and difficult, tasks ahead is defining harm. "Covered harm" under the bill is defined as any one of four types of harms, proximately caused by the use of a companion chatbot: (1) reasonably foreseeable physical or financial harm; (2) severe or reasonably foreseeable psychological or emotional harm to a reasonable child; (3) a highly offensive intrusion on a user's reasonable expectation of privacy; or (4) adverse discrimination against a user based on race, color, religion, national origin, disability, gender identity, sex, or sexual orientation.

First, is the drafting issue – noted above regarding the inconsistent use of the terms "child" and "user" – noting the difference between the reference to "child" in the second prong, general harm in the first prong, and harm to a "user" in the third and fourth prongs. "Child" should be used consistently throughout the bill.

Second, we seek more concrete definitions for the harms covered by the bill.

Three of the four harms identified in "covered harm" are somewhat subjective but largely actionable legal categories such as financial harm, privacy, and discrimination. These terms rely on more quantifiable or established legal standards that companies are familiar with and can navigate as they design and deploy products. We recommend specifically mentioning the applicable legal standards, such as "privacy rights protected by state or federal law" and "discrimination in violation of state and federal law."

The fourth harm, "severe and reasonably foreseeable psychological and emotional harm to a reasonable child" – particularly since the definition of "child" covers users in a wide of ages and developmental stages up to the age of 18, and some children use tools that are not designed specifically for minors-- is incredibly difficult to interpret, let alone implement. While the "reasonable child" standard is clearly intended to help, this language requires auditors, regulators, and courts, not to mention the developers and deployers, to determine how a hypothetical child of a similar age and developmental stage would react to a chatbot's output. Because children's emotional resilience and developmental stages vary widely, particularly if they have different needs or home support systems, what constitutes "severe" emotional harm can be highly subjective compared to other types of harm covered by the bill.

Audit regulations

The auditing provisions rise a host of concerns from the current insufficiency of the market for AI audits to the protection of sensitive business information to the broad grant of rulemaking authority to the Attorney General (including the authority of the AG to share audits with "qualified researchers"). We look forward to working with the authors to determine an appropriate way to promote accountability while protecting businesses' intellectual property, ability to innovate, and ability to offer products to users of all ages in California.

Third party audits generally do not advance safety objectives but, rather, impose a significant burden on the companies that must commission them. Audits must be rigidly conducted against clear standards, which are not generally suited to trying to evaluate a complex issue like the effects a platform has on minors.

Even academic researchers who have spent their entire careers studying the effects of technology use have not reached clear and convincing conclusions, given the impossibility of knowing the particular emotions and circumstances of an individual. This type of assessment would be completely outside the realm of expertise of any auditing firm.

Businesses subject to auditing requirements under European laws, such as the Digital Services Act, have incurred significant expense given a limited pool of qualified auditors, spent millions of dollars in extremal fees and countless personnel hours, and found themselves in a continuous loop of auditing where as soon as one audit ends, the next one must begin.

Given the importance of the issues at stake, and the dearth of successful AI audit regimes in other jurisdictions, the Legislature itself should define what role audits should play in the compliance framework for this legislation rather than simply leaving it to the AG to define through regulations. Nor should the bill grant the AG discretion to share company audit reports with any non-governmental third parties. While external researchers can play an important role in shaping academic research used in performing impact assessments required by Proposed Section 22612, sharing the sensitive information from business audit reports with "advocacy organizations" and "academic researchers" creates serious risks for businesses and may invite conflicts of interest without helping foster accountability.

Risk assessments

From a technical standpoint, Proposed Section 22612 requires an annual risk assessment of a companion chatbot that assesses five things: (1) the likelihood of covered harm occurring to users [effectively the likelihood of them violating the law subject to a private right of action]; (2) differential risks across age groups and developmental stages; (3) known vulnerabilities of children [unclear if this means known vulnerabilities of all children from a societal standpoint to the developer's knowledge or known vulnerabilities posed to children by the chatbot]; (4) empirical data from actual use [unclear how this is met if a risk assessment is being performed in advance of deploying a new chatbot]; and (5) relevant academic research and regulatory guidance. These requirements combined with the requirement that companies reasonably mitigate "any child safety risk" identified in a risk assessment, effectively seem to ask operators to foresee –and foreclose—against all future harms, which would not be possible.

Companies routinely conduct internal risk assessments and implement mitigation measures, their processes are generally flexible and context dependent. In contrast, the requirements imposed by **SB 1119** are overly prescriptive, product-specific, and recurring. In particular, the mandate to conduct annual, comprehensive risk assessments for a *specific* product feature, coupled with the detailed statutory criteria above, represents a level of granularity that would introduce operational and compliance complexities that merit careful consideration to ensure it is practicable, let alone effective.

Moreover, while internal risk assessments are a standard part of responsible product development, companies generally treat the underlying analyses as confidential, both to protect proprietary methods and to avoid legal exposure. Mandating public disclosure of detailed safety assessments for each companion chatbot would be highly atypical and could create operational, competitive, and liability challenges – disincentivizing candor and thoughtful assessments due to concern over those assessments becoming fodder for litigation.

Also, to the extent that existing chatbots are also subject to the risk assessment provisions, companies may need additional time to come into compliance.

Mandated warnings

We appreciate that this bill seeks to comprehensively address many of the important elements of a chatbot used by minors, including the product design, underlying data practices, user interactions, default settings, and parental controls. As several companies testified at the informational hearing of the Assembly Privacy and Consumer Protection Committee hearing last month, the industry is working diligently to solicit feedback from families about how they use these innovative tools, and to address families' concerns. We are working to gather specific feedback to the authors on these elements of the bill, including about product features being developed, the technical challenges of some of these proposals, and possible unintended consequences of some of these provisions. Additionally, we are working to evaluate how this bill's data restrictions would interact with other existing legal frameworks, most notably the CCPA so we can provide constructive feedback about the challenges companies will face in complying with this bill's proposals.

Prohibited conduct

As noted at the outset, we appreciate the general approach of allowing youth access to AI tools while ensuring that there are different levels of protection applied to their experiences through impact assessments, default protections, and parental controls. However, while well-intentioned, many of the provisions placing restrictions on outputs, default settings, crisis response expectations, and parental controls are not only detailed and prescriptive, but lack clear standards, making them difficult to interpret and operationalize in practice—particularly for dynamic, real-time conversational systems. By way of one example, Section 22612(d)(5) prohibits certain chatbot responses, such as attempting to "diagnose or treat" a child, as well as discouraging breaks or encouraging continued engagement. However, the bill does not clearly distinguish between harmful conduct and ordinary, good faith interactions, such as providing general wellness suggestions, or maintaining conversational engagement. As a result, operators may face uncertainty as to whether routine interactions fall within the scope of these prohibitions.

Taken together, these ambiguities create significant compliance challenges and potential liability exposure, even for operators acting in good faith to prevent bad outcomes. This uncertainty may ultimately discourage the deployment of tools altogether. Clarifying the boundaries between prohibited conduct and permissible interactions would improve implementation and better align the bill and its intended protective purpose.

Liability structure

Finally, we must raise concerns about **SB 1119's** excessively punitive liability structure. Under this bill, a public prosecutor can bring an action, even in the absence of harm, to recover undefined statutory penalties for any violation of this bill. They may even recover punitive damages without establishing harm.

While the private right of action does require actual harm for recovery, it too allows recovery of punitive damages. When combined with the incredibly vague mandates above, which make compliance particularly challenging and create expansive grounds for liability, this is particularly problematic.

Even more problematic, are the provisions providing for discrete violations. As a result, under Proposed Section 22616(c), each and every time a notice is not sent frequently enough throughout an interaction with a child (as "periodically" is not defined in the notice provision), it is an individual violation subject to a statutory fine of unknown dollars and punitive damage in an action by a public prosecutor; and also actual damages as well as punitive damages in a private right of action.

Additionally, each violation of any other provision – such as audits, prohibited behaviors, or data handling (each sale, each use, each sharing) -- is also its own discrete violation, subject to the same penalties, on an individual basis. This framework multiplies any liability, even minor mistakes like a missed disclosure, triggering multiple actionable violations subject to exceedingly high penalties. The operational challenge of monitoring every message and system obligation may discourage deployment tools designed specifically to help children play, learn, and grow.

New restrictions on selling, sharing, or using children's PI outside of the CCPA

Proposed Section 22613 states that an operator shall not "sell, share, or use for any purpose not expressly authorized by this chapter the personal information of a child." Note, this does not say "the personal information of a child collected by a companion chatbot." It is *all* PI of a child.

While the bill does not define the terms "sell", "share", "use" or "personal information," this effectively means that no operator under this bill can sell, share or use the personal information of anyone under the age of 18, unless expressly authorized by this bill (or future bill amending this chapter). Nothing in this bill, however, expressly authorizes the sale, sharing, or use, of PI of a child. And because "operator" under this bill includes any person who makes a companion chatbot available to user in this state—this applies equally to public entities and all businesses that may offer a companion chatbot to users (not just children) in this state.

While we hope to provide you with additional feedback, including language to address some of these concerns, for these reasons we must **OPPOSE UNLESS AMENDED SB 1119 (Padilla).**

Sincerely,

Ronak Daylami
Vice President for Advocacy | Privacy, Cybersecurity & Emerging Technologies
on behalf of

California Chamber of Commerce, Ronak Daylami
Computer & Communications Industry Association, Aodhan Downey
Civil justice Association of California, Annalee Augustine
Insights Association, Howard Fienberg
Software Information Industry Association, Abigail Wilson

cc: Legislative Affairs, Office of the Governor
Alexis Castro, Office of Senator Padilla
Consultant, Senate Judiciary Committee
Morgan Branch, Consultant, Senate Republican Caucus