June 16, 2026

TO: Members, Senate Judiciary Committee

SUBJECT: AB 2023 (WICKS, BAUER KAHAN) COMPANION CHATBOTS: CHILDREN'S SAFETY
OPPOSE UNLESS AMENDED – AS AMENDED APRIL 27, 2026
SCHEDULED FOR HEARING – JUNE 23, 2026

The California Chamber of Commerce and the undersigned respectfully must **OPPOSE UNLESS AMENDED AB 2023 (Wicks, Bauer Kahan)** as amended April 27, 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 provides the appropriate forum for stakeholders to balance these issues, rather than through ballot initiatives that risk limiting legislative and stakeholder input. 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. At the same time, technological innovation remains central to California's economic strength, its leadership in research and development, and its long-term social and economic progress, including benefits for children.

We appreciate **AB 2023**'s general approach of allowing youth access to AI tools while ensuring that different levels of protection are 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.

We also appreciate recent amendments removing provisions that would have permitted the Attorney General (AG) to share audit reports with independent child safety organizations or advocacy groups, narrowing third-party access to sensitive audit materials. However, other amendments have raised additional concerns regarding risk assessment requirements and their potential to imply guarantees against all harms. Nonetheless, we remain hopeful that suggestions below can help bridge the gap further to provide a workable framework for creating sensible safeguards around the development of chatbots used by children in this state while supporting continued innovation.

Defining "harm"

Perhaps one of the most important, and difficult, tasks ahead is defining harm. The prior iteration of this bill defined "covered harm" 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, under this prior version, the terms "child" and "user" were used inconsistently across this definition of "covered harm," resulting in ambiguity as to whether certain provisions were intended to apply only to minors or to users generally, with some prongs limited to children and others appearing to apply without an age restriction.

As we noted in our prior letter, “child” should be used consistently throughout the bill. While recent amendments largely replaced “user” with “child user” throughout the bill, it was noticeably not changed in several places and, in the definition of harm, the term “user” was simply deleted altogether—which arguably allows those elements of harm to be read more broadly as opposed to more narrowly, had they been limited to “child user” instead.

Second, we continue to urge more concrete definitions for the harms covered by the bill. As drafted, 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.” We appreciate that the amendments now tie privacy and discrimination harms to recognized legal standards. However, those clarifications do not resolve the separate issue regarding the bill's inconsistent references to "child" and "user." Both changes were necessary: one provides greater legal certainty regarding the nature of the harm, while the other ensures that the bill remains focused on harms experienced by children.

Our primary remaining concern remains centered around the inclusion of "severe and reasonably foreseeable psychological or emotional harm to a reasonable child” as a covered harm. Unlike financial harm, privacy violations, or unlawful discrimination, this category lacks clear, objective benchmarks that would allow developers, auditors, regulators, or courts to determine consistently when the standard has been met – 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.

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, maturity, 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.

Stated another way, from a practical standpoint, it will be very difficult for a developer to determine when an AI system's responses cross the line into conduct that creates liability under the statute. For example, the bill does not identify what distinguishes temporary distress, disappointment, embarrassment, anxiety, or loneliness—experiences that many children encounter in everyday life—from the type of severe psychological or emotional harm that would trigger liability under the statute. When combined with the significant liability exposure of this bill, smaller developers in particular are going to be in a precarious position if they do not know where the line falls. Ultimately, unlike the other covered harms, liability under this prong does not depend on a more readily identifiable event, such as a financial loss, privacy violation, or unlawful discriminatory act. Instead, it depends on predicting how a hypothetical child of a particular age and developmental stage would emotionally respond to a chatbot's output. As a result, developers may have difficulty determining in advance whether a given output creates liability under the statute.

We therefore encourage the authors to either remove this category of harm or provide additional objective criteria, examples, or limiting principles regarding what constitutes “severe and reasonably foreseeable psychological or emotional harm.” Without clearer standards, compliance obligations and potential liability will remain difficult to predict, particularly for smaller developers that lack extensive legal and risk-management resources.

Audit Requirements and Disclosure of Audit Information

AB 2023’s auditing provisions give rise to a host of concerns regarding unclear audit standards, a limited pool of qualified auditors, protections for sensitive business information, and broad AG disclosure authority,

including to “qualified researchers”.[1] Without more objective statutory standards, audit findings may vary significantly across auditors, reducing their reliability and utility for both regulators and operators.

While third-party audits can play a role in compliance frameworks, they also impose significant costs and may not reliably advance safety objectives for AI systems. Audits are most effective when conducted against clear, objective standards, which are not generally suited to trying to evaluate a complex issue like the effects a platform has on minors. Here, auditors may be asked to evaluate highly subjective concepts such as emotional harm, autonomy, decisionmaking, and excessively sycophantic behavior. These are inherently difficult assessments that remain the subject of ongoing academic debate. 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. In that context, these assessments are not readily susceptible to traditional audit methodologies and requiring third-party auditors to render definitive compliance judgments may produce inconsistent and unreliable results.

These challenges are compounded by the limited pool of qualified AI auditors and the significant costs associated with recurring audit obligations. For example, under European frameworks such as the Digital Services Act, companies have reported significant costs, limited auditor availability, and continuous audit cycles. Given the nascent state of the AI auditing sector and the lack of widely accepted auditing standards, the Legislature should carefully consider whether the audit regime currently included in the bill is likely to provide the intended regulatory value.

We are particularly concerned by **AB 2023’s** authorization for the AG to disclose "specific information" from audit materials to "qualified researchers." While we recognize the role of independent researchers in advancing child-safety research, the definition of a qualified researcher remains exceptionally broad. **AB 2023** provides limited guidance regarding the affiliations, qualifications, confidentiality obligations, security standards, or conflict-of-interest safeguards applicable to such recipients. We are especially concerned that these provisions could result in the disclosure of a company’s internal safety methodologies, testing protocols, or other sensitive operational information to third parties under the guise of “qualified researchers” given the incredibly broad definition of that term under the bill, which largely lacks meaningful safeguards to prevent access by entities with competitive interests adverse to the audited operator..[2] As a result, entities that satisfy the bill’s broad definition of a qualified researcher could potentially gain access

[1] While we appreciate recent amendments that remove prior language that would have permitted disclosure to advocacy organizations or child safety groups for the purpose of developing safety standards or educational resources, the bill continues to permit disclosure to “qualified researchers,” a category that remains broadly defined and subject to limited statutory safeguards, as further discussed in this section.

[2] As drafted, a qualified researcher would be defined as any individual or organization that does any of the following:

(1) Is **affiliated with** an academic institution, **nonprofit research organization, or independent research entity** or is otherwise **able to demonstrate relevant professional expertise**.

(2) **Demonstrates** a **legitimate** research purpose that is in the public interest and directly related to understanding, identifying, or mitigating risks to child safety or well-being arising from companion chatbots.

(3) Commits to conducting research in accordance with **applicable ethical standards** and is capable of complying with **applicable confidentiality, security, and data protection requirements.**

Notably, what these “applicable” security and data protection standards are; what the confidentiality requirements are; what types of affiliations qualify; how established a nonprofit or independent research entity has to be to be deemed legitimate enough to get this information; how one “demonstrates” they have a “legitimate” purpose – all of these things are left open ended, but these entities are being allowed to get **specific information from AI audits.**

to sensitive information regarding a company's safety practices, compliance approaches, or technical systems, creating risks that extend well beyond the bill's child-safety objectives.

As drafted, similar concerns arise with the requirement that the AG publish summaries of audit reports. Audit materials may contain information regarding internal testing methodologies, safety protocols, system architecture, risk assessments, mitigation measures, and other proprietary business information. Even where user information has been anonymized, disclosure of company-specific audit findings may create risks to intellectual property, information security, and fair competition. To the extent transparency is necessary, aggregated reporting regarding overall industry trends or high-level summaries may accomplish those goals while reducing the risk of inadvertently revealing information about a particular operator's systems, methodologies, or vulnerabilities. Aggregate reporting can promote accountability and transparency without the increased risk of exposing information that competitors could use to gain a commercial advantage or that malicious actors could use to identify weaknesses in a company's systems.

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 a role in shaping academic and child-safety research used in performing impact assessments, sharing the sensitive information from business audit reports with "academic researchers" or other outside entities without adequate safeguards and qualifications creates serious risks to intellectual property, information security, and fair competition among businesses while providing limited additional accountability benefits.

For these reasons, including the nascent stage of the AI auditing sector, we encourage the Legislature to consider a more targeted audit framework. Rather than imposing broad, recurring audit obligations across all covered operators, the bill could authorize the AG to request audit materials in connection with investigations of alleged noncompliance, require independent audits in circumstances where credible evidence of significant risk or repeated violations exists, and provide incentives for operators to conduct voluntary audits and implement best practices proactively. Such an approach would preserve regulatory oversight and accountability while better protecting confidential business information, encouraging innovation, and directing compliance resources toward situations where they are most needed.

AB 2023 should be amended to reflect that risk assessments are risk management tools, not guarantees against all future harm

Companies routinely conduct internal risk assessments, which are typically flexible, iterative, and context dependent. In contrast, the requirements imposed by **AB 2023** 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 that the framework is practicable, let alone effective.

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, an assessment of the likelihood of conduct giving rise to statutory liability under a private right of action, which under the current bill may be subject to disclosure to third parties beyond the AG]; (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", effectively seem to ask operators to foresee –and foreclose— against all future harms, which would not be possible. Recent amendments exacerbated these concerns, highlighting that developers must mitigate not only risks identified in the risk assessment but all risk.

Fundamentally, risk assessments are tools used to identify, prioritize, and mitigate risks — not guarantee that no harm will ever occur. We fully support the objective of protecting children from harm. However, effective risk-management frameworks recognize that risks vary in both likelihood and severity and therefore must be evaluated and addressed accordingly. No risk-management framework can function effectively if it assumes that every potential risk carries the same weight or demands the same level of mitigation. Responsible safety programs necessarily prioritize resources toward those risks most likely to cause significant harm.

A framework that effectively requires operators to mitigate every conceivable risk, regardless of magnitude or likelihood, risks creating compliance obligations that are both impracticable and less effective at advancing child safety. Stated another way, treating all risks as equivalent can divert attention and resources away from the most significant threats to children. The most effective approach is not to attempt to eliminate every conceivable risk, but rather to identify the most significant risks and implement mitigation measures proportionate to those risks.

AB 2023's requirement that operators reasonably mitigate "any child safety risk," when combined with a private right of action and broad liability standards, risks creating an expectation that operators must prevent all foreseeable harms before they occur. No product, service, or technology can be designed to eliminate all risk, particularly where outcomes depend on the individualized behavior, circumstances, and decisionmaking of users. A compliance framework should recognize that risk mitigation and risk elimination are distinct concepts.

This framework creates a significant risk of hindsight-based liability. Following an adverse event, it may be easy to identify risks that, in retrospect, appear foreseeable. The relevant question for operators, however, is what risks could reasonably have been identified and mitigated at the time the assessment was conducted based on the information then available. Without clear limiting principles, risk assessments may become less a tool for improving safety and more a mechanism through which operators are judged against information and outcomes that only became apparent after the fact.

These challenges are likely to be particularly acute for smaller developers who do not have the ability to dedicate substantial legal, compliance, policy, and technical resources to continually reassessing potential risks and documenting mitigation efforts against such broad and evolving standards. As a result, uncertainty regarding potential liability may discourage innovation or entry into the market altogether, even where developers are acting in good faith to implement reasonable safety measures. And while larger companies may be better positioned to absorb these costs, uncertainty regarding liability can affect operators of all sizes by diverting resources toward defensive compliance measures rather than safety innovation and product improvement.

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 – potentially disincentivizing candor and thoughtful assessments due to concern over those assessments becoming fodder for litigation.

For these reasons, we encourage amending **AB 2023** to clarify that risk assessments are intended to identify and mitigate reasonably foreseeable risks based on information available at the time the assessment is conducted, rather than serving as a basis for liability whenever a harm is later alleged to have been foreseeable. We also recommend allowing additional time for existing systems to come into compliance and only requiring subsequent risk assessments following a material change to a system, rather than on a fixed annual schedule regardless of whether meaningful system developments have occurred, in addition to clarifying and narrowing some of the elements of the risk assessment to address issues outlined above. Such revisions would better align the bill's requirements with established risk-management practices while reducing unnecessary compliance burdens, legal uncertainty, and barriers to innovation.

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 earlier in March, the industry is working diligently to solicit feedback from families regarding how they use these tools, to address their concerns, and to provide technical input to the authors on various elements of the bill, including product features under development, implementation challenges, potential unintended consequences, and interaction with existing legal frameworks such as the California Consumer Privacy Act (CCPA).

We do, however, note concerns with Proposed Section 22612(d)(4), which requires that an operator implement a mechanism for providing notice to a child user that the child is interacting with, or receiving content generated by, an AI system, by July 1, 2024. That notice must be reinforced periodically during extended interaction and must be presented in language and a format appropriate to a child.

While not objectionable in principle, its application is unclear, given the bill's definition of "child" as any individual under 18 years of age. This creates a compliance standard that must somehow simultaneously account for a wide range of developmental stages from early childhood through late adolescence (nearly adulthood), each with materially different levels of language comprehension and cognitive ability. As a result, the statute provides no clear guidance for compliance. A notice suitable for a teenager may be incomprehensible for younger users; while a notice for younger users may be overly simplistic or ineffective for older minors. This ambiguity creates uncertainty regarding how a compliant notice can be designed across such a broad age spectrum.

By contrast, tying the requirement to the intended user base of the chatbot, for example, or to defined age brackets or reasonable user segmentation, would provide a more clear and workable benchmark, improving implementation and better aligning the provision with its intended objective.

Prohibited conduct - clarity is needed to avoid compliance uncertainty and expansive liability

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 highly detailed and prescriptive, but they also lack clear standards, making them difficult to interpret and operationalize in practice—particularly for dynamic, real-time conversational systems. Among other things, we are concerned that the bill's rigid limitations on conversational memory or engagement presents complex issues, and hope to explore other options with the authors to balance safety concerns with operators' ability to provide positive user experiences.

By way of one example, Proposed Section 22612(d)(5) prohibits certain chatbot responses, such as attempting to "diagnose or treat" a child, unless the system is specifically designed and regulated for such purposes, and discouraging breaks or encouraging continued engagement. While well-intentioned, these prohibitions may, however, be difficult to apply to routine, low-risk interactions in systems designed to provide general support or conversational engagement. As drafted, 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.

For example, a student interacting with a school-based chatbot may express stress or anxiety about school. A response such as "you might try journaling, exercise, or mindfulness techniques to help manage stress" could be interpreted as general wellness guidance, but it is unclear whether it might be construed as attempting to "diagnose or treat" a condition under subsection (B). Similarly, a response such as "I'm here if you want to keep talking" could raise questions under subsection (F) regarding whether the system is improperly encouraging continued engagement or discouraging breaks.

Taken together, F. Clarifying the distinction between permissible general wellness guidance and supportive conversational engagement as opposed to prohibited conduct would improve implementation and better align the provision with its intended purpose.

Liability structure

Finally, we must raise concerns about **AB 2023's** excessively punitive liability structure. The bill permits public prosecutors to seek up to $5,000 per affected child per negligent violation, even absent harm, in addition to 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.

Particularly concerning are provisions that create discrete violations for each instance of noncompliance. As a result, under Proposed Section 22616(c), each 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 fines of up to $5,000-15,000 per affected child, 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 other provisions—such as audits, prohibited behaviors, or data handling (including each sale, each use, and each instance of sharing)—is also treated as a discrete violation subject to the same penalties on an individual basis. This structure effectively multiplies liability exposure, such that even minor or technical implementation errors can trigger multiple actionable violations subject to significant cumulative penalties.

As a result, for example, any failure to provide notice in an "age-appropriate format" or to repeat such notice "periodically" may create cascading liability. Because "periodically" is undefined, operators are left to guess the appropriate frequency of repetition in real time across interactions with users ranging from early childhood through late adolescence. If that judgment is later deemed insufficient, each subsequent message exchanged after the point at which notice should have been repeated may be treated as a separate violation, each independently subject to statutory penalties, punitive damages, and potentially actual damages in a private right of action.

The operational challenge of monitoring every message and system obligation may discourage deployment of tools designed specifically to help children play, learn, and grow. Taken together, the combination of vague performance standards, continuous real-time compliance obligations, and compounding per-instance liability creates a framework in which ordinary product design and operational decisions may be retrospectively converted into multiple statutory violations, increasing legal uncertainty and discouraging deployment of beneficial tools even for developers acting in good faith.

One way to begin to address such concerns would be to narrow the actionable violations to the affirmative obligations in Proposed Section 22612 which are most closely associated with risk of harm.

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" 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.

Sycophantic/excessively sycophantic

The bill uses to terms worth noting: “sycophantic”, defined to mean “validating of a user’s preferences or desires for the primary purpose or effect of optimizing engagement”; and, “excessively sycophantic”, defined to mean “sycophantic to an extent that is likely to have the substantial effect of subverting or impairing the user’s[3] autonomy, decisionmaking, or choice.”

First, independent of the policy considerations reflected within these definitions, we ask that you reconsider the use and definition of terminology that is clearly pejorative. Neutral language will better facilitate productive engagement and collaborative dialogue on issues of critical importance to Californians whereas this terminology risks contributing to unnecessary division. For example, *sycophantic* could have been more neutrally defined as “preference-driven engagement” and *excessively sycophantic* as “excessively preference-driven engagement.”

Second, terminology aside, as a matter of public policy, we urge consideration of an alternative framework that prohibits specific, identifiable conduct rather than relying on a standard that may be difficult for operators to interpret and apply consistently. While the bill defines "excessively sycophantic" conduct as behavior likely to have the substantial effect of subverting or impairing a user's autonomy, decisionmaking, or choice, it provides little guidance regarding how that threshold should be evaluated in practice. For example, many systems are intentionally designed to personalize content based on a user's stated interests, preferences, or requests in order to improve relevance and user experience. The bill does not clearly distinguish between permissible levels of personalization and conduct that would be deemed to impair a user's autonomy or decisionmaking.

Without more objective criteria, developers, auditors, regulators, and courts may reach inconsistent conclusions regarding when the statutory standard has been met. As a practical matter, uncertainty of this kind tends to favor larger market participants with more extensive compliance resources while making it more difficult for smaller developers to compete and innovate. That being said, even larger developers may struggle to implement and audit compliance programs when the underlying legal standard lacks clear, objective benchmarks, increasing the likelihood of inconsistent enforcement and defensive product design decisions that may not ultimately benefit users.

Continued concerns around scope

As noted in a previous letter, while **AB 2023** is largely framed as addressing concerns related to children's use of certain AI tools/ companion chatbots, several provisions within the bill continue to refer broadly to "users" rather than "children" or "child users." Because the bill defines "child" but does not define "user," these provisions may reasonably be interpreted to apply to all users, regardless of age. (The bill also fails to define “child user” but the inclusion of the term “child” at minimum has a narrowing effect to those individuals under 18 years of age.) As a result, the bill suggests that certain obligations applying to “users” are intended to extend beyond minors and apply to all uses, regardless of age, and not just those under eighteen.

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. Although recent amendments addressed many of these inconsistencies, others remain. For example, the bill continues to define an "operator" as a person that makes a companion chatbot available to a user in the state, rather than to a child user in the state. As drafted, this definition could encompass operators whose services are directed exclusively to adults, even though the stated focus of the bill is the protection of users under eighteen years of age.

If the Legislature intends for certain provisions to apply to adult users, the bill should clearly identify those provisions and distinguish them from requirements applicable only to children. Conversely, if the bill is intended to regulate companion chatbots only in the context of child users, ***all*** references to "user" should

[3] Note, it’s unclear why these references to “users” have not been modified to “child users” consistent with recent amendments.

be reviewed and revised to ensure the scope of the bill is clear and internally consistent. Clarifying the bill's scope would provide regulated entities with more fair and accurate notice of their obligations and reduce the risk of inconsistent interpretation or enforcement.

For these reasons we must **OPPOSE UNLESS AMENDED AB 2023 (Wicks, Bauer Kahan).**

Sincerely,

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

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

cc: Legislative Affairs, Office of the Governor
Samantha Huynh, Office of Assemblymember Wicks
Elise Gyore, Office of Assemblymember Bauer-Kahan
Consultant, Senate Judiciary Committee
Morgan Branch, Consultant, Senate Republican Caucus