

Evolving Expectations for AI in Mental Health



CONTRIBUTORS



Jodi Daniel



Andrea Linna



Eva F. Yin



Tracy Shapiro



Ty Kayam



Brandon Ge

CLIENT ADVISORIES

September 22, 2026

Artificial intelligence (AI)-enabled mental health products are becoming more ubiquitous with many individuals turning to chatbots, digital companions, coaching tools, and other AI systems for mental health support. They are also attracting increased attention from regulators, policymakers, and courts. As discussed in our previous client alert, *Legal Framework for AI in Mental Healthcare*, these products already fall within a patchwork of federal and state regulatory frameworks. Recent developments suggest government scrutiny of the sector is likely to continue increasing, including recent actions by agencies, state laws, and emerging litigation. For companies developing AI-enabled mental health products, key considerations may include medical device regulation, emerging state AI laws, Federal Trade Commission (FTC) oversight of health-related claims, privacy, transparency and health data use, and potential liability theories, including the unauthorized practice of medicine or psychology.

Mental Health AI Products and Regulatory Considerations

Mental health chatbots encompass a broad range of products, from tools that provide general information and wellness support to products that offer symptom assessment, wellness coaching, coping strategies, guided exercises, or other clinically oriented features. Some products are designed primarily to provide educational content, while others generate personalized responses, recommendations, or feedback based on user inputs. Although these products are often described as chatbots, AI companions, voice assistants, or other conversational tools, those labels generally describe how users interact with the technology rather than how it is regulated. Regulatory analysis typically focuses on the functions a product performs, the claims made about those functions, and the intended use of the product.

As products evolve to include additional functions, or if there are different claims or intended uses of existing functionalities, new regulatory issues may arise. For example, a mental health chatbot that initially provides general educational content may warrant additional analysis if it later incorporates features such as symptom assessment, screening, triage, or personalized recommendations.

Some products rely on predetermined decision trees or other deterministic approaches to generate outputs, while others use large language models to produce more open-ended outputs. The U.S. Food and Drug Administration's (FDA) August 2026 discussion paper, *Considerations for the Regulation of Generative AI-Enabled Medical Devices*, identifies areas for evaluating generative AI-enabled devices, including agentic devices. Please see our previous firm alert, *FDA Introduces a Potential Review Architecture for Generative AI-Enabled Medical Devices*, that explains this discussion paper in detail.

The manner in which a function is delivered can also be relevant. Similar functions may raise different considerations depending on the nature of the user interaction and the broader product experience. For example, a text-based chatbot, voice-enabled assistant, and AI companion may each



Nawa Lodin



Seamus Taylor

present mental health information in different ways, but each affects how users interact with that information.

FDA Regulation and Oversight

The main question that determines FDA oversight is whether a conversational AI product is a medical device. Under the Federal Food, Drug, and Cosmetic Act (FDCA), a product is a device if it is intended for use in the diagnosis of a disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease.¹ Consistent with its function-based approach to software regulation, the FDA evaluates intended use at the level of the software function rather than the product as a whole, looking at both design and claims.²

For many conversational AI products, a central regulatory question is whether the product is intended to support general wellness or a disease-related use.³ FDCA Section 520(o)(1)(B) excludes software functions intended to maintain or encourage a healthy lifestyle that are unrelated to the diagnosis, cure, mitigation, prevention, or treatment of a disease or condition.⁴ The FDA's revised [General Wellness Guidance](#), issued on January 6, 2026, also describes an enforcement discretion policy for low-risk general wellness products that may meet the device definition. The guidance includes general claims concerning relaxation, stress management, and mental acuity, and certain claims linking a healthy lifestyle to reducing the risk or impact of chronic diseases where the association is well understood and accepted. It distinguishes those claims from claims to diagnose or treat a specific disorder. Claims regarding disease-related uses, together with the product's functionality and other evidence of intended use, may support treatment of the product as a medical device and trigger further medical device analysis.

A separate pathway outside device regulation exists for certain clinical decision support (CDS) software, but it has important limitations for patient-facing products. Under FDCA Section 520(o)(1)(E), certain software intended to support healthcare professionals is excluded from the device definition.⁵ The FDA's revised [CDS Guidance](#), issued on January 29, 2026, clarifies the four statutory criteria for that exclusion. Most relevant here is the criterion that the function supports or provides recommendations to a healthcare professional and enables that professional to independently review their basis without relying primarily on the recommendations for an individual patient's diagnosis or treatment. A function intended to provide recommendations to patients or caregivers does not meet this criterion.

For products that do qualify as devices, the FDA has begun signaling how it expects developers to address the unique risks presented by generative AI. For example, the FDA's [November 6, 2025, meeting](#) of the Digital Health Advisory Committee focused on a hypothetical prescription generative AI-enabled therapy chatbot for adults with major depressive disorder. At the time of the meeting, the FDA had not authorized any generative-AI-based device for a mental health indication. Although the committee's recommendations are non-binding, the discussion provides insight into the issues the FDA and external advisors may consider relevant when evaluating these types of products. Specifically, the committee discussed the types of evidence that could support evaluation of such a product, including appropriate clinical comparators and outcome measures, as well as postmarket considerations such as crisis escalation, overreliance and excessive use, performance drift, and use outside the product's intended scope. The committee also discussed product labeling, including disclosures regarding the AI's purpose and limitations, when users should seek human care, and privacy-related considerations. These topics were discussed as part of the hypothetical case study and should not be viewed as establishing new regulatory requirements or formal FDA policy, but they do provide some insight into issues the FDA may take into consideration.

State Law Developments

In parallel with FDA regulatory efforts, and in the absence of a broader federal regulatory framework, a growing number of states have enacted laws targeting the use of AI technologies in healthcare, with some specifically targeting AI in the delivery of mental healthcare.

Several states, including California ([SB 243](#)), Georgia ([SB 540](#)), New York ([SB 3008](#)), Oregon ([SB 1546](#)), Rhode Island ([S 2195/H 7350](#)), and Utah ([HB 452](#)), require companion chatbot or mental health chatbot operators to clearly and conspicuously disclose that users are interacting with AI rather than a human, with some jurisdictions imposing recurring disclosure obligations (e.g., New York requires disclosure at the outset of an interaction and at least every three hours thereafter; Utah requires disclosure after any seven-day lapse in use). A parallel and increasingly common requirement is the implementation of crisis-response protocols: California, Georgia, New York, Oregon, and Rhode Island

all mandate that operators maintain mechanisms to detect and address expressions of suicidal ideation or self-harm and refer users to crisis services.

A second cluster of state laws focuses on scope-of-practice restrictions rather than disclosure. Colorado (HB 1195), Illinois (HB 1806), Maine (HP 1397), Nevada (AB 406), Rhode Island (S 2197/H 7349), Tennessee (SB 1580), and Vermont (H 816) generally prohibit AI from independently making therapeutic decisions, engaging directly in therapeutic communication, or generating unreviewed treatment recommendations. Instead, these states limit the use of AI to administrative and supplementary support functions (e.g., note-taking, transcription) subject to licensed-professional oversight and, in some states, client consent. New Mexico imposes an informed-consent requirement specifically for counselors and therapists using AI for tasks such as diagnosis or treatment planning. Texas takes a somewhat different approach, permitting AI-assisted diagnostic recommendations subject to practitioner review and disclosure obligations. California is currently considering SB 903, which would extend similar administrative-support-only limitations to psychotherapy specifically, along with consent requirements for AI-assisted recording, transcription, or triage.

In addition, several states have enacted laws establishing more comprehensive regulatory frameworks for the use of AI in a variety of contexts, including mental health.⁶ Similar to the state laws described above, such frameworks require disclosures to consumers (and in some cases the ability to opt-out) when AI technologies are used to assist with or make certain decisions for such consumers. Last, states are actively considering and drafting additional mental health AI legislation, and we expect to see a growing number of bills introduced and considered in state legislatures in the coming years. Companies operating in this space should expect the patchwork to continue expanding and should evaluate compliance on a state-by-state basis as new laws take effect.

Health Privacy Considerations

Mental health conversational AI products raise a distinct set of privacy considerations on top of the FDA and state AI-specific frameworks discussed above. The Health Insurance Portability and Accountability Act and its implementing regulations (HIPAA) generally would not apply to a direct-to-consumer chatbot, even if it collects highly sensitive mental health information, unless it is provided for or on behalf of a HIPAA-covered entity such as a health plan or a healthcare provider that bills health plans. Where a tool implicates federally assisted substance use disorder treatment, federal regulations at 42 C.F.R. Part 2 governing the confidentiality of substance use disorder records (Part 2) could apply, though this use case is rarer.

To the extent that HIPAA and/or Part 2 does not apply, and depending on the states where users reside, state consumer health privacy laws such as Washington's My Health My Data Act (MHMDA) could apply. These laws generally require opt-in consent for the collection and sharing of consumer health data and could, in the case of the MHMDA, create significant litigation exposure due to its private right of action. In addition, many state comprehensive consumer privacy laws classify information revealing a mental health diagnosis as sensitive health data, which could also trigger heightened requirements such as opt-in consent for processing.

FTC Considerations

Section 5 of the FTC Act, which prohibits "unfair or deceptive acts or practices in or affecting commerce," is an important source of federal consumer protection exposure for AI mental health products. Under the FTC Act, a representation is generally "deceptive" if it is material and likely to mislead a reasonable consumer, and a practice is "unfair" if it causes or is likely to cause substantial injury that consumers cannot reasonably avoid and that is not outweighed by countervailing benefits to consumers or competition. AI mental health tools could trigger FTC oversight based on unsubstantiated or false claims of efficacy or clinical capability, undisclosed use of AI or failure to disclose that users are talking to a bot, and deceptive claims about privacy practices.

The FTC also enforces the Health Breach Notification Rule (HBNR), which is broad in scope and applies to many non-HIPAA-covered health apps. Notably, the HBNR treats any unauthorized disclosure of identifiable health information as a reportable breach and is not limited to just cybersecurity intrusions or nefarious behavior. Thus, even an unauthorized disclosure to a vendor could constitute a breach under the HBNR, triggering requirements to notify affected users, the FTC, and, for larger breaches, media outlets.

The Children's Online Privacy Protection Act (COPPA) generally could apply where an AI chatbot is directed to children (which is assessed based on the totality of the circumstances) or the operator has actual knowledge it is collecting personal information from a child under 13. Among other obligations,

COPPA requires verifiable parental consent and a compliant privacy notice before collecting personal information from children under 13. Guidance also clarifies that COPPA requires parental consent to disclose children's personal information to third parties to train or otherwise develop AI technologies. In general, the impact of AI chatbots on children appears to be an area of focus for the FTC, which issued [orders](#) in September 2025 to seven companies that operate consumer-facing AI-powered chatbots seeking information on how the companies measure, test, and monitor potentially negative impacts of the technology on children and teens.

State Law Liability Considerations

The use of AI products in the delivery of mental healthcare implicates a variety of considerations under state law for the companies and providers who deploy them, including questions related to the unlicensed practice of a profession, corporate practice of medicine, and professional liability. At the state level, there is ongoing litigation involving consumer-facing chatbots in which plaintiffs have asserted theories of product liability, negligence, wrongful death, and consumer protection claims.

Many state laws require that, in order to deliver certain mental health services (e.g., the delivery of therapeutic interventions), individuals must hold licenses issued by a state board or agency (e.g., the state board of psychology). For example, with certain exceptions, California law prohibits the practice of psychology without a license issued by the California Board of Psychology.⁷ A person who practices psychology without a license is guilty of a misdemeanor punishable by imprisonment in the county jail not exceeding six months, or by a fine not exceeding two thousand dollars (\$2,000), or by both.⁸ Separately, in states recognizing the corporate practice of medicine doctrine, the doctrine may apply to arrangements involving AI products. California Business and Professions Code Section 2400 generally denies corporations and other artificial legal entities professional rights, privileges, and powers under the Medical Practice Act, subject to statutory exceptions, including for qualifying professional medical corporations.⁹

The application of these licensure and corporate practice frameworks to AI products used in the delivery of mental health services is an evolving area, and the treatment of any particular product or use case will depend on its specific facts and circumstances, including how the product is designed, marketed, and deployed, and the role of licensed professionals in connection with its use.

Design and Governance

The following design and governance practices, while not exhaustive and necessarily shaped by how a product is developed and deployed, may help mitigate risk and are generally easiest to incorporate before launch.

- *Crisis escalation:* State laws in California and New York require that companies that deploy AI companion chatbots must maintain crisis prevention protocols to prevent the chatbot from producing suicidal ideation, suicide, or self-harm content to the user and refer the user to crisis service providers when necessary.¹⁰
- *Scope controls and safeguards:* Given the tendency of generative AI-enabled conversational AI to shift based on the intent of the user, consider refusal mechanisms, retrieval constraints, and limits on clinical outputs, including testing controls under realistic conditions and across supported modalities and analyzing patterns from repeated interactions (such as sycophancy).
- *Human oversight:* While state requirements differ, meaningful human oversight is particularly important. Some states prohibit licensed professionals from allowing AI to engage directly in therapeutic treatment, while others address professional review and approval of therapeutic recommendations or treatment plans.
- *Postmarket monitoring:* Ongoing monitoring can help identify drift, bias, and other safety issues that emerge after deployment, and postmarket monitoring is of particular significance to the FDA. The FDA's November 2025 advisory committee discussed monitoring for performance drift, adverse events, overuse, and changes in the scope of use while the August 2026 discussion paper separately seeks input on whether, and under what conditions, greater reliance on postmarket monitoring could address residual uncertainty at authorization.
- *Clear labeling that the AI is not a human therapist:* Patient-facing products utilizing conversational AI should incorporate clear and regular notifications to users (in compliance with applicable law) that they are interacting with AI technology and not a human therapist. Companies should also evaluate whether their products are subject to applicable state AI laws imposing disclosure, transparency, or other consumer protection requirements.
- *Disclaimers:* The product should include comprehensive disclaimers tailored to product functions. If a product will only provide informational or wellness content, disclaimers should include that the company does not provide medical advice and it is not a substitute for the advice

of a licensed healthcare provider. If the product has not been reviewed by the FDA or another regulatory authority for use in a medical or clinical context, disclaimers should acknowledge this as well. California Health and Safety Code Section 1339.75 prescribes placement for disclosures in covered written, audio, and video patient communications and requires instructions for contacting an appropriate human, subject to the human-review exception described above. Note, however, that disclaimers alone do not determine FDA oversight, which may further consider labeling, advertising, design, and other relevant evidence.

Outlook

The regulatory landscape for mental health conversational AI is evolving through FDA oversight, state AI and professional practice laws, privacy and consumer protection requirements, and litigation under existing liability theories.

Importantly, regulatory questions regarding mental health conversational AI remain unresolved. The FDA has not yet issued specific guidance on these types of products, state laws continue to expand, and key liability theories are still being tested in court. For instance, most litigation to date concerns allegations and procedural rulings, rather than established liability rules for all mental health conversational AI products. Companies should expect continued change and consider regulatory issues early in development and frequently post-deployment.

Contact Us

Wilson Sonsini works with clients from across the digital health ecosystem, including developers, vendors, and deployers of mental health AI products. For more information, please contact Wilson Sonsini attorneys [Jodi Daniel](#), [Andrea Linna](#), [Eva Yin](#), [Tracy Shapiro](#), [Ty Kayam](#), [Brandon Ge](#), [Nawa Lodin](#), [Seamus Taylor](#), or any member of Wilson Sonsini's [Digital Health practice](#) or [Data, Privacy, and Cybersecurity practice](#).

[1] 21 U.S.C. § 321(h).

[2] 21 U.S.C. § 321(h); 21 C.F.R. § 801.4.

[3] FDCA § 520(o)(1)(B); 21 U.S.C. § 360j(o)(1)(B).

[4] 21 U.S.C. § 360j(o)(1)(B).

[5] 21 U.S.C. § 360j(o)(1)(E).

[6] *See, e.g.*, 11 Cal. Code Regs. § 7000 et seq.; Colorado SB 26-189 (2026).

[7] California law defines the “practice of psychology” as rendering or offering to render to individuals, groups, organizations, or the public any psychological service involving the application of psychological principles, methods, and procedures of understanding, predicting, and influencing behavior, such as the principles pertaining to learning, perception, motivation, emotions, and interpersonal relationships; and the methods and procedures of interviewing, counseling, psychotherapy, behavior modification, and hypnosis; and of constructing, administering, and interpreting tests of mental abilities, aptitudes, interests, attitudes, personality characteristics, emotions, and motivations. Cal. Bus. & Prof. Code § 2903(a).

[8] Cal. Bus. & Prof. Code § 2970.

[9] Cal. Bus. & Prof. Code §§ 2400, 2402, 2408.

[10] *See* Cal. SB 243; N.Y. Gen. Bus. Law Article 47.