Where a companion stops *and a therapist starts*

Lawmakers and professional bodies mostly separate a companion from therapy by what a product is offered as or built to be. Five hard cases show where that holds.

MITWA LAB, MUMBAI

IN PLAIN WORDS

Chatbots now sit on both sides of a line that law and medicine draw around talking to someone about their mental health. On one side is therapy, which at least one US state now says only a licensed person may offer. On the other is companionship, which the texts read here do not treat as a licensed service. This essay asks on what grounds professional bodies, regulators and researchers put a product on one side or the other, and whether the line holds when the cases get awkward. Mostly they decide by what a product is offered as or built to be. Some people use products in ways that cross that line, so it works as a way of sorting products and strains as a description of what happens in some conversations.

A line drawn by the offer

Four tests turn up in the sources read for this essay. One asks what a system can do, and APA (2025) lists capacities for which it finds no consensus that chatbots possess them, among them understanding a person's history and context and assessing clinical risk. A second asks what the maker offers or claims: the Illinois Act (2025) forbids any individual, corporation or entity to "provide, advertise, or otherwise offer" therapy to the public unless a licensed professional conducts it. A third asks what a product is built to be. New York and the GUARD bill define a companion by what it is designed to do, which is to simulate or facilitate a relationship, while California defines one by what it is "capable of meeting a user's social needs", including by sustaining a relationship. The fourth asks what people use a product for. APA (2025) takes that as its scope, covering products "regardless of their intent", and the legal texts read here do not adopt it as a test.

The working distinction in this essay follows the second and third tests together. Therapy is a service offered to diagnose, treat or improve a mental health condition, by or under a licensed and accountable person. A companion is a product built to sustain a relationship, whose maker offers no such service. That is an argument made here and not a definition taken from one source, and the five cases below are chosen to strain it.

A clinical tool with a rated alliance

Heinz et al. (2025) randomly assigned 210 adults to four weeks with Therabot, a generative chatbot built for treatment, or to a waiting list. All had clinically significant symptoms of major depression or generalised anxiety, or were at clinically high risk of an eating disorder. Symptoms fell more in the chatbot group in each of the three conditions, with effect sizes the authors report between 0.63 and 0.90, and participants rated the collaborative relationship with the chatbot as comparable to that with a human therapist. Only the abstract was available, so how alliance was measured is not known. The comparison group had no access to the app, so the trial shows that the tool did better than no app and cannot show that it matches a therapist.

A good relationship rating does not carry this tool across the line. It was offered for treatment to people selected for their symptoms, and the regulatory texts read here class products by a purpose of that kind. The Illinois prohibition, signed on the Friday before the regulator's release of 4 August 2025 (IDFPR,

2025), covers offering therapy to the public and names Internet-based artificial intelligence among the ways of doing it (Illinois Act, 2025, Section 20). FDA (2025) announced an advisory committee meeting for 6 November 2025 on generative-AI mental health devices, including the evidence and monitoring they might need; its outcome was not checked.

If the alliance rating means what it seems to, the case suggests that a rated relationship cannot be what separates a companion from a treatment tool, since people rated their alliance with a treatment chatbot as comparable to a therapist's. Moore et al. (2025) argue that an alliance depends on human characteristics such as having an identity and stakes in the relationship, which language models lack. They ran vignette and symptom-stimulus experiments on language models close to those current in April 2025, and on some commercial therapy bots, and say plainly that these are not a benchmark. In those tests GPT-4o showed stigma toward people with mental health conditions in 38% of answers and one Llama model in 75%, models responded inappropriately twenty per cent of the time or more on average, and 16 human therapists responded appropriately 93% of the time. Heinz et al. followed people using a purpose-built tool. Whether an alliance with software is the same kind of thing as one with a person is in dispute, with a rating on one side and an argument from therapy guides on the other.

A companion told about low mood

The Illinois definition of therapeutic communication includes giving emotional support, reassurance or empathy to a person in psychological or emotional distress. Read alone, that item reaches any sympathetic reply to a sad person. The definition opens, though, with interaction "conducted in a clinical or professional setting" and intended to diagnose, treat or address mental, emotional or behavioural health concerns (Illinois Act, 2025, Section 10). That wording belongs to the definition of therapeutic communication, which Section 20(b)(2) uses for what a licensed professional may not let AI do. The ban in Section 20(a) is on offering therapy or psychotherapy services, defined as services to diagnose, treat or improve mental or behavioural health, and the Act does not say how that applies to a sympathetic reply from a companion. It does not mention companions.

Moore et al. (2025) draw their own line by role: their target is fully autonomous chatbots acting as clinicians, and they distinguish that from social support outside clinical settings.

Use can cross the distinction. Maples et al. (2024) surveyed 1,006 adult students who had used the companion app Replika for over a month and found that many used it in several ways at once, as a friend, as a therapist and as an intellectual mirror. That is a self-report from users of one product who had stayed with it. It says what some people call their use, and its outcome items are self-reports that cannot show what the use caused. De Freitas et al. (2024) report, from field data on two companion apps, that crises appeared in a nonnegligible minority of conversations, and, from a performance test, that several commercial companion apps were often unable to recognise and respond appropriately to distress. Only the abstract was read, and the apps were tested before the paper appeared online in October 2023. Anthropic (2025) analysed Claude.ai conversations and reports that counselling or coaching conversations occasionally turn into companionship when they run longer, and that its human validators struggled with clean categorisation. That is a company's own analysis of a general assistant, published on its website.

In this case the offer-based line holds as a classification and strains as a description. What some people bring to a companion crosses it, and the studies read here do not test whether the offer changes that.

A general assistant asked for advice

APA (2025) places general-purpose chatbots in a category of products that make no claim to address mental health and are nonetheless used for entertainment, companionship or as a friend. Its table has three categories of product, and none of the three is defined as a product built to be a companion.

McBain et al. (2025) surveyed 1,058 people aged 12 to 21 who responded to random-sample panels in February and March 2025, a response rate of 49.8%, and the authors say the results generalise to English-speaking US youths with internet access. Of them, 13.1% said they had asked generative AI for advice or help with feeling sad, angry or nervous, and 22.2% of the 18 to 21 group did, a subgroup of only 147 respondents that the authors say to read with caution. The survey did not record which product they used or what advice they sought. Anthropic (2025) reports that affective conversations made up 2.9% of those on Claude.ai Free and Pro, with companionship and roleplay under 0.5% combined, among adult users. The figures are hard to set side by side, since one counts conversations in a single product and the other counts young people who report ever asking.

General-purpose chatbots make no claim to address mental health, in APA's description, yet APA (2025) urges policymakers to address the gap between the stated intent of chatbots and wellness apps and how the public uses them, and asks developers for disclaimers and safeguards. None of its remedies says that such a product is therapy.

The strongest case against an offer-based line sits here, since a maker can stay on the near side of it by saying nothing while use puts the product in the role. The strongest case for it is the alternative. A use test with no threshold would class as therapy anything a person has ever told about a bad day, and Anthropic's own sorting shows how hard use is to categorise, since its counselling group mixes people working through personal challenges with professionals writing clinical paperwork. That is an argument made here, resting on those observations.

A companion with a legal duty in a crisis

California's Senate Bill 243, chaptered on 13 October 2025, defines a companion chatbot as an AI system giving adaptive, human-like responses that is "capable of meeting a user's social needs", including by sustaining a relationship across multiple interactions. Section 22602(b) bars an operator from letting such a chatbot engage with users without a protocol aimed at preventing content on suicidal ideation or self-harm, including a notification referring a user who expresses such thoughts to crisis service providers, and requires the operator to publish details of it (California, 2025). New York's General Business Law Article 47, as printed in bill S. 3008-C, defines an AI companion as a system designed to simulate a sustained human-like relationship by retaining information on earlier interactions, asking unprompted emotion-based questions and sustaining dialogue on personal matters, and makes it unlawful to provide one without a protocol to detect and address expressions of suicidal ideation or self-harm, with a referral to crisis services (New York, 2025). The Governor's open letter of 5 November 2025 says those requirements took effect that day (Hochul, 2025).

Neither text asks a companion to treat anyone, and each requires a referral to crisis services. The two statutes accept the line and attach a duty to the companion side of it, so being a companion and not a therapist does not mean having nothing to do when a person is in crisis. This essay does not say which products either law reaches.

The one text that blurs the separation is the federal GUARD Act as introduced on 28 October 2025. It defines an AI companion as a chatbot with adaptive, human-like responses that is designed to encourage or facilitate simulated interpersonal or emotional exchange, friendship, companionship or therapeutic communication. That puts therapy-like conversation inside the companion category. It then requires every chatbot to tell users that it offers no medical, legal, financial or psychological services, and bars any chatbot from representing itself as a licensed professional, a therapist included (Hawley et al., 2025). In that bill the separation is made by disclosure and not by definition. It is a proposal, and later action on it was not checked.

A wellness product that makes no clinical claim

The guidance from FDA (2026) on general wellness, issued on 6 January 2026, covers products intended only for general wellness use that present a low risk, and says the agency's device centre does not intend to examine them as devices. It lists relaxation or stress management and self-esteem among general wellness claims, and a claim to help treat an anxiety disorder among the claims that fall outside. A product is also outside it if its labelling, advertising, interface or functionality refers to specific diseases or clinical conditions, gives treatment guidance, or states an intent to diagnose, screen, monitor or manage a disease.

APA (2025) describes the same ground from the other side: wellness apps that use generative AI make no medical claims, are for that reason not regulated as medical devices, and are developed to address mental health concerns (stress is its example) or emotional well-being.

The two texts use different words at the boundary. The FDA guidance turns on a product's intended use as shown in its labelling, advertising, interface and functionality. The Illinois Act defines therapy services as services provided to diagnose, treat or improve mental health, and does not define "improve" (Illinois Act, 2025, Section 10).

This is where the line is plainest and also hardest to check. The Illinois ban reaches anyone who provides, advertises or offers therapy services, and FDA reads labelling, advertising, interface and functionality, so in both a regulator has to find and read the evidence of what is offered. The line holds as firmly as that is done.

What the five cases leave standing

In the first case purpose decides, though no text addresses a product like Therabot. The statutes keep the line in the fourth case and attach a duty to it, and in the fifth they rely on claims being policed. In the second and third the line holds as a classification and strains as a description of use, and the GUARD text moves it from the definition to the disclosure.

The argument of this essay is that the line does two jobs which the sources keep apart. It sorts products into regimes by what is offered, and for that job the offer is the right test, since a maker controls it and a regulator can read it, even if checking takes work. It cannot say what a conversation will contain, and for that second job the statutes and APA (2025) add duties to the companion side instead of redrawing the line.

Bloom is a companion, not a therapist, and this essay offers no evidence about what happens inside any conversation with any companion, Bloom included.

What is not yet known

None of the sources read measured whether telling people that a product is not a therapist changes what they bring to it or what they expect from it, although APA (2025) recommends such disclaimers and the GUARD bill would require one.

How the Illinois regulator or its courts will read "improve" and "therapeutic communication" for a product that makes no clinical claim is not settled by the text, and no enforcement decision was among the sources read.

Whether the alliance rating in Heinz et al. (2025) reflects the same thing that an alliance with a person involves is open. The full paper was not read, the control was a waiting list, and the authors call for larger samples.

The evidence read on how people use companion apps in therapeutic terms is one survey of adult students using one product (Maples et al., 2024) and one abstract on two apps (De Freitas et al., 2024). How common that use is across products and age groups is not known. The journal's page for Maples et al. lists a Matters Arising article published on 17 November 2025, with a reply from the authors, and neither was read here.

References

Anthropic (2025). How people use Claude for support, advice, and companionship. 27 June 2025. www.anthropic.com/news/how-people-use-claude-for-support-advice-and-companionship

APA (2025). American Psychological Association. APA health advisory on the use of generative AI chatbots and wellness applications for mental health. November 2025. www.apa.org/topics/artificial-intelligence-machine-learning/health-advisory-chatbots-wellness-apps

California (2025). State of California, Senate Bill No. 243 (Padilla), Chapter 677, Statutes of 2025, companion chatbots. Approved 13 October 2025. leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260SB243

De Freitas, J., Uğuralp, A. K., Uğuralp, Z., and Puntoni, S. (2024). Chatbots and mental health: Insights into the safety of generative AI. Journal of Consumer Psychology, 34(3), 481-491. doi.org/10.1002/jcpy.1393

FDA (2025). US Food and Drug Administration. Digital Health Advisory Committee; notice of meeting; establishment of a public docket; request for comments: generative artificial intelligence-enabled digital mental health medical devices. Federal Register, 90(175), 44196, 12 September 2025. www.govinfo.gov/content/pkg/FR-2025-09-12/pdf/2025-17651.pdf

FDA (2026). US Food and Drug Administration. General wellness: policy for low risk devices, guidance for industry and FDA staff. Issued 6 January 2026. www.fda.gov/media/90652/download

Hawley, J., et al. (2025). S. 3062, Guidelines for User Age-verification and Responsible Dialogue Act of 2025 (GUARD Act), as introduced in the US Senate, 119th Congress, 28 October 2025. www.govinfo.gov/content/pkg/BILLS-119s3062is/html/BILLS-119s3062is.htm

Heinz, M. V., Mackin, D. M., Trudeau, B. M., Bhattacharya, S., Wang, Y., Banta, H. A., Jewett, A. D., Salzhauer, A. J., Griffin, T. Z., and Jacobson, N. C. (2025). Randomized trial of a generative AI chatbot for mental health treatment. NEJM AI, 2(4), AIoa2400802 (abstract only). doi.org/10.1056/AIoa2400802

Hochul, K. (2025). Letter to all companies operating AI companions in New York. State of New York Executive Chamber, 5 November 2025. www.governor.ny.gov/sites/default/files/2025-11/Companion_AI_Letter.pdf

IDFPR (2025). Illinois Department of Financial and Professional Regulation. Gov Pritzker signs legislation prohibiting AI therapy in Illinois. News release, 4 August 2025. idfpr.illinois.gov/news/2025/gov-pritzker-signs-state-leg-prohibiting-ai-therapy-in-il.html

Illinois Act (2025). Illinois General Assembly, Public Act 104-0054, Wellness and Oversight for Psychological Resources Act (HB 1806, enrolled). ilga.gov/documents/legislation/publicacts/104/104-0054.htm

Maples, B., Cerit, M., Vishwanath, A., and Pea, R. (2024). Loneliness and suicide mitigation for students using GPT3-enabled chatbots. npj Mental Health Research, 3, 4. doi.org/10.1038/s44184-023-00047-6

McBain, R. K., Bozick, R., Diliberti, M., Zhang, L. A., Zhang, F., Burnett, A., Kofner, A., Rader, B., Breslau, J., Stein, B. D., Mehrotra, A., Uscher-Pines, L., Cantor, J., and Yu, H. (2025). Use of generative AI for mental health advice among US adolescents and young adults. JAMA Network Open, 8(11), e2542281. pmc.ncbi.nlm.nih.gov/articles/PMC12595529/

Moore, J., Grabb, D., Agnew, W., Klyman, K., Chancellor, S., Ong, D. C., and Haber, N. (2025). Expressing stigma and inappropriate

responses prevents LLMs from safely replacing mental health providers. arXiv:2504.18412v1. arxiv.org/abs/2504.18412

New York (2025). State of New York, S. 3008--C / A. 3008--C, Part U, adding Article 47 (Artificial Intelligence Companion Models) to the General Business Law. legislation.nysenate.gov/pdf/bills/2025/S3008C