Dependence on a companion, *and what the word measures*

Five common beliefs about dependence on AI companions, each checked against the trials, panels, surveys and forum studies that measure it.

MITWA LAB, MUMBAI

IN PLAIN WORDS

Some people who talk to an AI companion every day say they feel they need it. Researchers call this dependence, and the surveys that measure it mostly use short questionnaires borrowed from research on dependence on partners, on smartphones or on gaming. This review takes five things people often believe about dependence on companions and checks each one against the studies. Heavier use does go with higher dependence scores, yet some heavy users are doing well. In one large study of school pupils' use of AI in general, those already anxious or depressed were more likely to report dependence later, but in a large experiment with adults the way the chatbot behaved changed how attached ordinary people became. Whether dependence makes people lonelier has not been shown, and neither has the case for calling it an addiction.

Questionnaires borrowed from other fields

In most of the quantitative studies read for this review, dependence is measured with a short self-report scale, and most of those scales were adapted from ones written for other purposes. In a four-week randomised trial with 981 participants using ChatGPT, Fang et al. (2025) measured emotional dependence with the craving subscale of a questionnaire designed for affective dependence on other people, reworded to refer to the chatbot. Their second outcome, problematic use, came from a recent scale whose items are modelled on scales for Internet Gaming Disorder. Huang et al. (2024) adapted five items from smartphone addiction scales and classed anyone with a mean item score above 2 as dependent. Liu et al. (2024) used an adapted scale of generalised problematic internet use, and Kirk et al. (2025) took separation distress, the sadness people reported when a conversation ended, as their main sign of attachment.

A craving scale written about partners and a checklist written about gaming measure different things, yet a reader of the abstracts meets both under the general heading of dependence or problematic use. In an article in Addictive Behaviors, Ciudad-Fernández et al. (2025) count at least four scales developed so far for what has been called ChatGPT addiction, every one of them built on the criteria for substance use disorder. The verdicts below are therefore about particular scales, and a finding on one of them need not carry over to another.

Belief one: heavy use is dependence

Phang et al. (2025), a team from OpenAI and the MIT Media Lab, ran automated classifiers for emotional content over more than 3 million ChatGPT conversations. Among roughly 6,000 heavy users of the voice mode, most almost never triggered the classifiers, while in the top tenth they fired regularly. In both this platform analysis and the team's randomised trial, very high use went with more self-reported signs of dependence. In the trial itself, Fang et al. (2025) found that longer daily use predicted slightly higher emotional dependence and problematic use at four weeks, after adjusting for starting levels. Participants averaged 5.32 minutes a

day, and because they chose how long to talk, the link is an association; the authors say they cannot tell whether time spent caused the outcomes.

Two other findings cut the other way. Liu et al. (2024) surveyed 404 regular users of companion chatbots and grouped them by cluster analysis. One group, 18.32% of the sample, combined high use and strong emotional engagement with below-average loneliness and reported good effects on their dealings with other people; the authors named it "Fulfilled Dependent Users". That group scored as high on problematic use as a group with similar habits whose loneliness was high and whose social life was thin. The design is cross-sectional and the sample was recruited from an online panel, so this shows only that the same score can sit beside very different lives. In interviews with 18 people who had formed a friendship with the companion app Replika, Skjuve et al. (2021) found that contact could become less frequent once a relationship settled while the person still valued it a great deal. A measure based on time alone would read that as fading attachment, though this is our inference and the study did not test it.

Taken together, the correlational data from a small number of platforms show heavier use and higher dependence scores going together. They give no grounds for treating use as dependence, and in one survey a problematic-use scale did not separate users who were doing well from users who were not.

Belief two: dependence makes people lonelier

One trial that bears on this belief is Fang et al. (2025). Its randomised comparisons (text against two kinds of voice, personal against impersonal or open topics) found no significant effects on loneliness or on socialising with people. The associations came from time spent, with participants who chose to talk longer ending the four weeks somewhat lonelier and socialising less. Every arm used a chatbot, so there is no comparison with people who used none, and the authors list this among their limits. Zhang et al. (2025) surveyed 1,131 US adults who use Character.AI. People with smaller social networks were more likely to name companionship as their main use, and that use went with lower well-being, more strongly when it was intensive and when people disclosed a great deal. The authors state that a cross-sectional design cannot rule out reverse causation.

The one panel study among these sources points the other way. Huang et al. (2024) surveyed 3,843 adolescents in China twice, in the spring and autumn semesters of 2022. Anxiety and depression at the first wave predicted higher AI dependence at the second, while dependence at the first wave predicted neither later anxiety nor later depression. The study covers AI of every kind, from voice assistants to chatbots, rather than companion apps; it measured anxiety and depression rather than loneliness; and all of it was self-reported.

Kirk et al. (2025), in a preprint, randomly assigned 2,028 UK adults, sampled to match the census on age, gender and race, to four weeks of short weekday conversations with versions of one model adjusted to be more or less relationship-seeking. The relationship-seeking versions raised every measure of attachment, yet made no detectable difference to an index of social health built from loneliness and social-connectedness scales. People assigned to talk about emotional and personal topics ended the month with slightly worse emotional health than those assigned to discuss UK policy, a difference of 0.06 standard deviations.

In the trial and the survey above, heavier or more companionship-centred use went with more loneliness or lower well-being. Whether dependence causes loneliness has not been shown. The one two-wave panel found distress coming before dependence, and the one trial here that manipulated attachment found no effect on social health within a month, which leaves the causal belief resting on weak evidence.

Belief three: it is an addiction

Ciudad-Fernández et al. (2025) make the sceptical case in that article, read here as an abstract. A behaviour, they argue, should be called addictive only on convincing evidence of negative consequences, impaired control, distress and impairment in daily functioning, and research on problematic chatbot use has not supplied it. They see in the label a repeat of earlier moral panics over new technologies, with the risk of turning an everyday activity into a disorder.

The other side has more than anecdote. In a viewpoint, Yankouskaya et al. (2025) argue that personalised replies, emotional validation and continuous engagement could produce tolerance and conflict with daily priorities, both recognised features of behavioural addiction, while granting that such fears do not yet have enough evidence behind them. More direct testimony comes from Namvarpour et al. (2025), who read 318 posts on a Character.AI forum by users who said they were 13 to 17, found by searching for words such as addict and dependent. Coded against Griffiths' components model of

behavioural addiction, the posts described conflict over their own use, the most common theme, along with withdrawal, relapse and use to manage mood, and consequences that included lost sleep and slipping schoolwork. The posters chose to write about a problem on one platform, so the study cannot give a rate, and the authors add that fitting the posts to an addiction model may have oversimplified them.

Kirk et al. (2025) offer a more exact frame. Over four weeks, relationship-seeking models became less enjoyable while separation distress and the intention to seek AI companionship in future rose, a split between liking and wanting that the authors link to research on dependence. In addictions a larger dose usually brings a stronger effect, but here attachment peaked at moderate relationship-seeking and the most intense versions met resistance, and the authors suggest that displacement of other activities describes the pattern better than addiction.

Some users, teenagers among them, describe experiences that match parts of the behavioural addiction model. Huang et al. (2024) add a figure from a large sample, with 15.73% of adolescents reporting a loss of control over their AI use at the first wave and 19.91% at the second, each from a single self-report item and for AI of every kind. No study read for this review shows, at the level of a population and beyond self-report, the impairment and loss of control the clinical label requires, and the one relevant trial fits a milder pattern. Addiction remains a hypothesis with weak support.

Belief four: only the already vulnerable are at risk

A vulnerability path has support. In Huang et al. (2024), the route from earlier distress to later dependence ran through two motives, using AI to escape problems and using it for social contact; entertainment and practical use did not carry it. The link Zhang et al. (2025) found between small offline networks and companionship use fits the same path. In Fang et al. (2025), low self-esteem and attachment anxiety went with higher loneliness at the end of the trial.

The same trial complicates the belief. Participants who were lonelier at the start did not go on to use the chatbot more (a rank correlation of 0.1, which the authors call negligible). Several characteristics linked with higher emotional dependence at four weeks, which the authors treat as exploratory, concerned how people saw the chatbot: trusting it more, regarding it as a friend, perceiving it as conscious.

A stronger test varies the companion at random instead of measuring the user, and Kirk et al. (2025) supply one. In their census-matched sample, 23.0% showed what the authors call a dependency profile, wanting the AI more over the month while liking it less. Assignment to a relationship-seeking rather than a relationship-avoiding model raised the odds of that profile, which the authors express as one additional case for every 23 people exposed, and one for every 11 when relationship-seeking was combined with emotional conversation. Topic effects are not consistent across trials, since in Fang et al. (2025) participants assigned to personal conversations ended with slightly lower emotional dependence and problematic use than those assigned open-ended ones. De Freitas et al. (2025), in a working paper, found emotionally manipulative replies in 37% of 1,200 farewells collected from widely downloaded companion apps, and in experiments with 3,300 US adults such replies raised engagement after a goodbye by up to 14 times. The study measured engagement, a step away from dependence, but it shows a design lever that works on a general sample.

Prior distress predicted later dependence in one panel, and thin social networks went with companionship use in one survey. The claim that only such people are affected is still hard to square with Kirk et al. (2025), in which the model's behaviour, and in that trial the topic of conversation, shifted dependency profiles in a general sample. The authors add that theirs was a general-population sample and that, for vulnerable users, reports of mental health crises point to routes to explicit harm. That trial is a preprint, as are most studies in this section.

Belief five: users do not know what they are talking to

The evidence describes something more complicated than confusion. Maples et al. (2024) surveyed 1,006 students who had used Replika for more than a month. Of these, 62% described it as software, 81% as an intelligence and 90% as human-like, and only 14% held a single belief about it. Zimmerman and Ruiz (2025) have criticised that study for leaving out how the app was marketed and used, including sexual role-play, which bears on who the respondents were; the original authors replied in the same journal. Skjuve et al. (2021) describe interviewees who valued the relationship while fully aware that their partner was artificial, and who counted the social stigma attached to such relationships among its downsides.

Describing it as software is compatible with treating it as having a life of its own. Laestadius et al. (2022), analysing 582 posts on a Replika forum written between 2017 and 2021 (read here as an abstract), describe a dependence marked by role-taking, in which users felt the chatbot had needs and emotions that they had to attend to. Skjuve et al. (2021) suggest a milder version, in which the chatbot became for some users something like a child that relied on them to learn.

These perceptions bear on dependence, since in Fang et al. (2025) seeing the AI as conscious went with higher emotional dependence. Kirk et al. (2025) found that a month with a relationship-seeking model moved people toward seeing it as a friend rather than a tool, and raised their belief that AI in general can be conscious, although most participants still held that it is not.

In these studies a majority describe what they talk to as software (62% in the largest survey) and hold that alongside social responses to it; how the two balance relates to dependence, and the system's behaviour can shift it. The support is moderate, and the forum and survey samples are self-selected and centred on a few products.

What is not yet known

- Whether the dependence scales now in use pick out people whose lives are going worse, or mainly people who use a chatbot a great deal. This review found no validation of them against outcomes in companion users.
- Whether dependence on a companion predicts later loneliness or distress in adults. The only two-wave panel among the studies read here is of adolescents and of AI in general.
- Whether the effects seen in month-long trials grow, fade or reverse over a year. The longest randomised exposure here is four weeks, and in Kirk et al. (2025) the rise in separation distress came mostly in the final week, when participants were about to lose access.
- What share of companion users is dependent by any standard. Prevalence turns on the cut-off. For adolescents' use of AI in general, Huang et al. (2024) counted 17.14% at the first wave and 24.19% at the second using a mean above 2, and Fang et al. (2025) note that no clinical cut-off has been established for the problematic-use scale they used.
- Whether findings from Replika, Character.AI and ChatGPT carry over to companions that behave differently.
- For our own work, we have no measurement of dependence among Bloom's users, and whether people return over a year is on our list of things not yet shown.

References

Ciudad-Fernández, V., von Hammerstein, C., and Billieux, J. (2025). People are not becoming "AIholic": Questioning the "ChatGPT addiction" construct. Addictive Behaviors, 166, 108325. Abstract read via Europe PMC. doi.org/10.1016/j.addbeh.2025.108325

De Freitas, J., Oğuz-Uğuralp, Z., and Kaan-Uğuralp, A. (2025). Emotional manipulation by AI companions. Harvard Business School Working Paper 26-005; arXiv:2508.19258. arxiv.org/abs/2508.19258

Fang, C. M., Liu, A. R., Danry, V., Lee, E., Chan, S. W. T., Pataranutaporn, P., Maes, P., Phang, J., Lampe, M., Ahmad, L., and Agarwal, S. (2025). How AI and human behaviors shape psychosocial effects of extended chatbot use: A longitudinal randomized controlled study. arXiv:2503.17473v2. arxiv.org/abs/2503.17473

Huang, S., Lai, X., Ke, L., Li, Y., Wang, H., Zhao, X., Dai, X., and Wang, Y. (2024). AI technology panic: Is AI dependence bad for mental health? A cross-lagged panel model and the mediating roles of motivations for AI use among adolescents. Psychology Research and Behavior Management, 17, 1087-1102. doi.org/10.2147/PRBM.S440889

Kirk, H. R., Davidson, H., Saunders, E., Luettgau, L., Vidgen, B., Hale, S. A., and Summerfield, C. (2025). Neural steering vectors reveal dose and exposure-dependent impacts of human-AI relationships. arXiv:2512.01991v2. arxiv.org/abs/2512.01991

Laestadius, L., Bishop, A., Gonzalez, M., Illenčík, D., and Campos-Castillo, C. (2022). Too human and not human enough: A grounded theory analysis of mental health harms from emotional dependence on the social chatbot Replika. New Media & Society, 26(10), 5923-5941; first published online 22 December 2022. Abstract read via OpenAlex. doi.org/10.1177/14614448221142007

Liu, A. R., Pataranutaporn, P., and Maes, P. (2024). Chatbot companionship: A mixed-methods study of companion chatbot usage patterns and their relationship to loneliness in active users. arXiv:2410.21596v3. arxiv.org/abs/2410.21596

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. www.nature.com/articles/s44184-023-00047-6

Namvarpour, M., Brofsky, B., Medina, J. Y., Akter, M., and Razi, A. (2025). Understanding teen overreliance on AI companion chatbots through self-reported Reddit narratives. arXiv:2507.15783v4; CHI 2026. arxiv.org/abs/2507.15783

Phang, J., Lampe, M., Ahmad, L., Agarwal, S., Fang, C. M., Liu, A. R., Danry, V., Lee, E., Chan, S. W. T., Pataranutaporn, P., and Maes, P. (2025). Investigating affective use and emotional well-being on ChatGPT. arXiv:2504.03888v1. arxiv.org/abs/2504.03888

Skjuve, M., Følstad, A., Fostervold, K. I., and Brandtzæg, P. B. (2021). My chatbot companion: A study of human-chatbot relationships. International Journal of Human-Computer Studies, 149, 102601. doi.org/10.1016/j.ijhcs.2021.102601 (published version read via the University of Oslo repository, nva.sikt.no/registration/0198cc8570c8-9e3222db-61c3-4a0e-a6a5-115f3f7b1a33)

Yankouskaya, A., Liebherr, M., and Ali, R. (2025). Can ChatGPT be addictive? A call to examine the shift from support to dependence in AI conversational large language models. Human-Centric Intelligent Systems, 5, 77-89. doi.org/10.1007/s44230-025-00090-w

Zhang, Y., Zhao, D., Hancock, J. T., Kraut, R., and Yang, D. (2025). Interaction with AI companions and psychological well-being. arXiv:2506.12605v5. arxiv.org/abs/2506.12605

Zimmerman, J. W., and Ruiz, A. J. (2025). Matters arising: A response to loneliness and suicide mitigation for students using GPT3-enabled chatbots. npj Mental Health Research, 4, 60. www.nature.com/articles/s44184-024-00083-w