

Advertising to the distressed: The commodification of mental health data in AI chatbots

Posted on April 14, 2026

By Nicole Gross and Hannah van Kolfshooten



Generative Artificial Intelligence (genAI) chatbots have become an important outlet for many people around the world who are experiencing mental health issues. Of its 800 million weekly users, around 10 percent use ChatGPT for [emotional support](#) while more than one million use the chatbot to talk about issues such as depression, [psychosis and suicidal ideation](#). For many users, these conversations feel intimate, tailored and supportive, like talking to a good friend or therapist. Yet these interactions take place within commercial digital platforms whose business models rely on collecting, analysing and commodifying user data. As Big Tech companies explore the use of behavioural advertising in genAI systems, influenced by the global market for health data being worth an estimated [\\$74.41 bn](#), profound ethical tensions arise between the commodification of user data and the protection of vulnerable individuals seeking mental health support.

On one hand, the World Health Organization states that mental health is a [fundamental human right](#), which means that everyone has the right to available, accessible, acceptable, and good-quality mental healthcare. As more than [1 billion people](#) live with mental health issues and many do not receive adequate care, more and more turn to genAI chatbots, which are often free, accessible and highly interactive. As AI chatbots are becoming quasi-therapeutic spaces, they exert an influence over the right to mental health. However, unlike traditional mental healthcare contexts where confidentiality and professional duties restrict the use of patient information, AI chatbots operate within a commercial ecosystem that is designed to maximise engagement and monetise user attention. This means that users pay a high yet often invisible price for their mental health.

The commodification of mental health data works as follows: tech companies collect and analyse any information that appear in prompts, chat logs and other user interactions. While some chatbots, such as [Grok](#), claim that “we do not sell your personal information or use it for marketing”, their privacy policies clearly show that aggregated and de-identified data is being ‘shared’ with affiliates and third parties, including advertisers and other business partners (Table 1). Grok started to run ads in [August 2025](#) and ChatGPT has been trialling ads since [February 2026](#).

Table 1 – Privacy Policy Extracts from ChatGPT, Gemini, Grok, and CharacterAI

Chatbot	Extracts from Privacy Policies	Link
ChatGPT	“To assist us in meeting business operations needs and to perform certain services and functions, we disclose Personal Data to vendors and service providers, including providers of hosting services, customer service vendors, cloud services, content delivery services, support and safety services, email communication software, web analytics services, payment and transaction processors, search and shopping providers, and information technology providers.”	https://openai.com/policies/row-privacy-policy/
Google Gemini	“We may share non-personally identifiable information publicly and with our partners – such as publishers, advertisers, developers or rights holders.”	https://policies.google.com/privacy
Grok	“We may use your personal information for a variety of purposes...” “For example to develop new product features...to operate and expand our business activities, to identify new customers, and for data analysis.”	https://x.ai/legal/privacy-policy
CharacterAI	“We use information for the following purposes:”... “Provide advertising and recruit new users, including provision of tailored advertising.””	https://policies.character.ai/privacy

This type of [psychographic profiling](#) is highly problematic: by analysing prompts and chatlogs, genAI systems can infer sensitive information about a person’s psychological state and emotional vulnerabilities. The ethical concern is that such signals can then be used to refine advertising strategies that target individuals precisely at moments when they are most vulnerable. In practice, this means that a user who confides to a chatbot that they feel anxious or lonely may unknowingly generate a behavioural profile that categorises them as a “stress reactor,” “easily deflated,” or “receptive to emotional messaging.” A person who turns to a chatbot to talk about insomnia or anxiety may later be exposed to advertisements for sleep supplements or expensive therapy apps, for instance. Or someone discussing body image could be categorised as responsive to aspirational messaging and subsequently shown ads for cosmetic procedures, weight-loss programmes or confidence-boosting services. In these cases, deeply personal disclosures become inputs into a commercial system that predicts and commodifies behaviour.

While users are warned in (admittedly lengthy and complex) T&Cs that they should not include personal or sensitive information in their prompts, many people remain unaware just how much data genAI platforms collect and where their data goes, and this raises serious concerns about manipulation, exploitation and the right to privacy. Arguably, people could always choose not to use chatbots or read the T&Cs carefully and set their privacy controls accordingly. However, this is an arduous process as users are typically opted-in by default and platforms may still personalise ads even when ad personalisation is [turned off](#). Platform power and information asymmetries reinforce Big Tech’s business models while rendering users choiceless over what happens their data once it enters the commercial ecosystem. At the same time, the ads users are exposed to may even exacerbate the very mental health struggles that led them to seek support in the first place.

Most genAI systems are [neither strictly risk-assessed nor cohesively regulated](#) globally. While the EU’s regulatory frameworks – GDPR, the AI Act and the Digital Services Act – apply to genAI systems, none of these ‘gold standard’ regulations directly address the risks that arise when AI chatbots use personal and emotional disclosures in an advertising-driven ecosystem. The EU’s recent proposal of a [Digital Omnibus](#) is also further shift away from regulation and towards the endorsement of competitiveness. Without stronger safeguards, the commodification of emotionally vulnerable interactions risks turning moments of psychological distress into opportunities for commercial exploitation by chatbots.

Authors/Affiliations:

Nicole Gross, Associate Professor in Business & Society, National College of Ireland

Hannah van Kolfshooten, Postdoctoral Research Fellow, University of Basel

Conflicts of Interest: None to declare

(Visited 12 times, 12 visits today)

[Artificial intelligence](#), [Mental Health](#)

[« PREVIOUS POST](#)

0 Comments

Login ▼

LOG IN WITH

OR SIGN UP WITH DISQUS

Share

[Best](#) [Newest](#) [Oldest](#)

Be the first to comment.

[Subscribe](#) [Privacy](#) [Do Not Sell My Data](#)

BMJ Blogs

Comment and Opinion | Open Debate

The views and opinions expressed on this site are those of the original authors. They do not necessarily represent the views of BMJ and should not be used to replace medical advice. Please see our full Blog [Terms and Conditions](#).

All BMJ blog posts are posted under a CC-BY-NC licence

[BMJ Journals](#)

SEARCH

CATEGORIES

Select Category ▼

FORUM EDITORS

[Owen Schaefer](#) and [Kathryn Muyskens](#)

BMJ CAREERS

