Case Reports and Reviews

Review article

Drafting Distress in a Chat Window: A Review of AI-Mediated Help-Seeking And Care Pathways

Dr Zoe Wyatt

Wyatt Potage Consulting, Australia

Correspondence

Dr Zoe Wyatt

Wyatt Potage Consulting, Australia

- Received Date: 18 Feb 2026
- Accepted Date: 24 Feb 2026
- Publication Date: 26 Feb 2026

Keywords

mental health, help-seeking, conversational AI, chatbots, self-disclosure, stigma, referral, digital therapeutics

Abstract

Conversational AI is increasingly used to rehearse and articulate distress before people seek formal mental health support. This review examines evidence across three clinically relevant domains: self-disclosure to conversational systems, stigma-related barriers that shape uptake, and the conditions under which AI use supports movement from private coping toward real-world care pathways. Findings suggest disclosure in chatbot contexts is sensitive to perceived anonymity, privacy expectations, trust cues, and fear of judgement, and may provide short-term emotional relief. Self-stigma and label avoidance are associated with attitudes toward AI-delivered support, potentially lowering the threshold for initial engagement while also increasing the risk of avoidance maintenance. Early feasibility research indicates that screening and referral chatbots can be acceptable and may reduce navigation burden, and outcome evidence shows that some structured conversational interventions can reduce distress in specific populations. The review consolidates these results into a mechanism-led synthesis of where AI use is most likely to assist early help-seeking, where it can stall, and the bridge conditions that increase the probability of transfer to human or formal supports.

Introduction

Generative AI is increasingly being used as a first-line way to think through distress, especially when the alternative is slow, expensive, or feels too exposing in the moment [1]. Researchers describe this shift in plain terms: conversational tools are being used for mental health-related support because they are always available, feel private, and can be accessed without the social friction that often comes with telling another person what is going on [1]. That matters clinically because the first steps of help-seeking are often the hardest steps. Many individuals do not begin by contacting a professional [2]. They begin by trying to make sense of what is happening, working out whether it represents a mental health difficulty, and gauging whether disclosure is worth the interpersonal and emotional cost [2]. Contemporary help-seeking models describe this as a staged process that includes symptom recognition and appraisal, perceived need, and stigma-related barriers before formal support is sought [3]. In other words, the early phase is not "treatment-seeking," it is meaning-making and risk assessment, and that is exactly where conversational AI is starting to sit.

This is also where the disclosure literature becomes relevant. Conversational AI is likely to shape help-seeking not because it is magically therapeutic, but because it changes the conditions under which people practise language for distress [4]. Experimental research shows that perceived anonymity and reduced fear of judgment can increase willingness to disclose intimate information to a chatbot, even when people still report greater trust in humans for sensitive topics [4]. A recent review further maps how privacy expectations, social presence, and conversational features influence disclosure to conversational AI, reinforcing that disclosure is not a fixed trait but a context-sensitive behaviour [5]. When you place those findings alongside usage data, the relevance becomes harder to dismiss as a US trend. In a nationally representative US survey, 13% of 12–21-year-olds reported using generative AI for mental health advice [6], while Australian survey data indicate that 9.9% of adults reported asking ChatGPT health-related questions within the previous six months [7]. Australian mental health-specific work adds that both community members and clinicians are already encountering AI tools as a form of emotional support and self-guided coping, alongside mixed views about benefit and risk

Citation: Wyatt Z. Drafting Distress in a Chat Window: A Review of AI-Mediated Help-Seeking And Care Pathways. Case Rep Rev. 2026;6(1):090.

[8]. Finally, the international footprint is visible in mixed-methods research drawing participants from 29 countries who repeatedly used ChatGPT for emotional and mental health support [9].

Mapping this landscape is not about positioning AI as a substitute for therapy. It is about recognising that conversational AI is already appearing early in the help-seeking sequence, shaping how distress is articulated, how disclosure is rehearsed, and whether a person chooses to involve a human. This article synthesises current evidence across three linked domains: disclosure processes in chatbot interactions, the role of self-stigma and anticipated judgement in uptake and engagement, and the conditions under which AI use appears to support movement toward formal care rather than reinforcing a private coping loop. The paper concludes by identifying practice-relevant implications and specific gaps that warrant further empirical investigation.

Before reviewing each evidence domain in detail, the next section consolidates the literature into a mechanism-led view of where AI use is most likely to help or stall during early help-seeking, and identifies leverage points that keep use oriented toward onward action.

Failure modes and leverage points in AI-mediated early help-seeking

Failure mode 1: relief without movement

A consistent finding across disclosure research is that disclosure can be experienced as emotionally relieving, including in chatbot contexts [4]. Relief is not trivial in mental health terms because it can reduce arousal and create momentary coherence. The failure mode emerges when relief becomes the endpoint. If the interaction repeatedly produces short-term emotional settling without any change in the person's appraisal, behaviour, or access steps, the pattern can become self-reinforcing. Longitudinal work on human–chatbot relationships supports the plausibility of this dynamic by showing that disclosure patterns can shift over time, including evidence that conversational breadth may narrow as use becomes more established [13]. The risk is not "chatting is bad", it is that repeated narration can become a loop that regulates affect while leaving the help-seeking pathway unchanged.

Leverage point: convert relief into a single decision output. The most mechanism-consistent shift is to move from narration to a decision prompt that forces translation into one concrete next step, such as a shareable summary for a clinician or a single booking or referral action [3,10]. This keeps the reinforcing quality of relief, while linking it to follow-through.

Failure mode 2: the trust ceiling

Experimental work suggests a useful asymmetry. Perceived anonymity and reduced fear of judgement can make disclosure feel easier with a chatbot, even while people report greater trust in a human interlocutor for sensitive topics [4]. Reviews of disclosure to conversational AI reinforce that disclosure is context-sensitive and shaped by privacy expectations, social presence, and conversational features [5,12]. The failure mode here is a trust ceiling. People may share enough to feel heard or clarified, but hesitate to go further, hesitate to act on guidance, or hesitate to take the output into a human relationship. Trust concerns do not have to be explicit to shape behaviour. They can operate as a quiet cap on depth, accuracy checking, and willingness to transfer the narrative to a clinician.

Leverage point: produce a human-compatible artefact, not more disclosure. The mechanism-consistent move is not to push the person to share more into the same channel, but to translate what has been said into a concise, clinically usable summary that can be carried into care. This also aligns with the review-level message that disclosure is not a trait, it is shaped by context, and contexts change when a narrative becomes shareable and owned by the person rather than held inside a private loop [5,12].

Failure mode 3: stigma bypass that becomes avoidance maintenance

Self-stigma is repeatedly linked to attitudes toward help-seeking and appears associated with willingness to consider AI-delivered psychotherapy in young adults [11]. Related findings suggest that some users perceive chatbot support as more discreet or less constrained by stigma-related barriers than traditional services, even while remaining uncertain about benefit [14]. Label avoidance adds an additional mechanism. When a formal help-seeking step is experienced as adopting a stigmatised identity, a low-exposure channel can feel tolerable precisely because it avoids interpersonal visibility [10]. The failure mode is when the tool functions as a stigma bypass that maintains avoidance. The person can externalise distress without confronting the interpersonal step that would reduce avoidance long-term, leaving the underlying identity threat intact.

Leverage point: planned hand-off that preserves low exposure while increasing contact. The most defensible leverage point is to use the interaction to design a staged transfer, moving from private to interpersonal disclosure in a deliberately bounded way, rather than assuming that comfort in a chatbot will spontaneously generalise into human help-seeking. This sits cleanly within what the evidence can support: stigma-related processes shape what feels tolerable, and tolerability influences uptake and engagement [10,11,14].

Failure mode 4: navigation stall at the action threshold

Help-seeking models emphasise that movement from appraisal to service contact often breaks down at the point where distress must be translated into action [3]. Feasibility work on screening-and-referral chatbots is relevant here because these systems are designed to reduce navigation friction and guide users through a structured sequence that culminates in recommendations [10]. The failure mode is not lack of insight, but decision load: uncertainty about what level of support is indicated, the cognitive burden of choosing under stress, and the practical complexity of navigating options.

Leverage point: a structured pathway output with a single external action. The evidence most directly supports a pathway-oriented use of conversational systems when the interaction ends with a recommendation that is linked to an acceptable next step, rather than a general conversation that can continue indefinitely [3,10]. This keeps the clinical focus on movement, not on the chatbot as the endpoint.

Summary proposition

Across these mechanisms, the most clinically useful synthesis is that conversational AI appears most likely to function as early scaffolding when it produces transferable outputs that lower exposure and navigation burden while increasing the probability of interpersonal contact or structured follow-through. It is most likely to drift into a closed loop when the primary reinforcement is emotional relief without action,

Figure 1. *Mechanism-led model of AI-mediated early help-seeking*

Help-seeking pinch point	What AI can support	How it can stall	Bridge condition
Symptom articulation	Put vague distress into workable language.	Relief without movement: repeated narration, no clearer plan.	A clinician-ready summary: symptoms, duration, triggers, impact, what tried, what next.
Disclosure safety	Test wording before interpersonal exposure.	Rehearsal becomes avoidance: no transfer to a person or service.	Two shareable drafts: short opener and fuller version for a message, email, or first appointment.
Action threshold	Reduce decision friction for the next step.	Decision paralysis: repeated 'what do I do' querying.	One completed external step: booking request sent, call made, referral submitted, or appointment time secured.
Navigation burden	Organise logistics while access is delayed.	Indefinite waiting: planning continues, no escalation rule.	A time-limited waiting plan plus an escalation trigger: what changes means seek professional help now.
Stigma and label avoidance	Lower-exposure entry point when disclosure feels too exposing.	Stigma bypass becomes maintenance: private support is the endpoint.	A staged hand-off: who to tell first, exact wording, specific request, and a date to do it.

when trust concerns cap depth and transfer, or when stigma-related avoidance is bypassed without a planned hand-off [4,10,11,13,14].

Disclosure and the psychology of telling

Self-disclosure to conversational AI has become an empirical topic because disclosure is not a neutral act in mental health support, it is a behaviour shaped by perceived costs and safety [12]. Experimental evidence indicates that disclosure in chatbot contexts is associated with specific perceptions, particularly anonymity, trust, and fear of judgement [12]. In a controlled study comparing disclosure to a chatbot versus a human, Croes and colleagues examined willingness to disclose intimate information and tested underlying processes including perceived anonymity and trust, showing that these perceptions help explain why disclosure can feel easier in chatbot interactions, even where trust is higher for a human interlocutor [4]. Their work also links disclosure experiences with emotional relief, which provides a grounded explanation for why a chatbot interaction may be experienced as subjectively helpful at the point of disclosure, regardless of what happens next [4].

Longitudinal evidence adds an important correction to one-off "do people disclose to bots" debates by showing that disclosure can develop over time within human–chatbot relationships [13]. Skjuve and colleagues examined self-disclosure across time and explicitly describe disclosure in terms of breadth and depth, noting that disclosure patterns may change as the relationship matures, including evidence that conversational breadth may be reduced with maturation [13]. This is the empirically defensible point: disclosure to chatbots is not static, and ongoing interaction can be associated with shifts in what is shared and how the interaction is used [15]. Skjuve, et al. [13] do not link these disclosure shifts to clinical outcomes; rather, they show that disclosure patterns can change over time with repeated chatbot use, consistent with the development of relationship-like interaction features.

A broader synthesis helps explain why disclosure findings can appear inconsistent across studies and systems. Papneja and Yadav [12] propose an emergent framework in which self-disclosure to conversational AI is shaped by interacting determinants, including interface modality and conversational features, user characteristics, mediating mechanisms such as perceived anonymity and social presence, and contextual factors including privacy expectations. This matters because it supports a precise conclusion that is consistent with the evidence base: chatbots do not uniformly increase disclosure, but disclosure is sensitive to the interaction context and to the mechanisms identified in empirical and synthesis work. Finally, disclosure cannot be treated as culturally uniform [4,12,13,16]. Chin and colleagues analysed large-scale chatbot interactions and reported cultural differences in depressive mood expression across Western and Eastern countries, demonstrating that what is expressed and how it is expressed in chatbot contexts varies across cultural groupings [16]. Taken together, these studies support a bounded claim: conversational AI can change the conditions under which disclosure occurs, and those conditions include perceived anonymity, trust-related perceptions, privacy expectations, and culturally patterned expression.

Self-stigma and the low-exposure first step

Self-stigma functions less like a background attitude and more like a cognitive filter that shapes how distress is interpreted and whether help-seeking feels permissible [17,18]. Clinically, it is often observed through avoidance, minimisation, delay, and a preference for managing difficulties privately until they feel more coherent, certain, or socially defensible [19]. Self-stigma functions as an identity threat: disclosure and help-seeking can be appraised as signalling weakness, inadequacy, or personal failure, which in turn predicts delay and avoidance in moving toward support [20-22].

In practical terms, this increases the perceived cost of help-seeking and shifts the balance toward concealment [11]. It also helps explain why low-exposure forms of support can

feel appealing at the earliest stage [4] (Croes et al., 2024). Conversational AI can function as an on-demand channel for articulating distress with reduced immediate interpersonal exposure, and experimental evidence indicates that perceived anonymity and lower fear of judgement are associated with greater willingness to disclose in chatbot interactions [4]. This does not imply that chatbots reduce stigma as a construct; rather, it supports a behavioural interpretation that lower-exposure channels may lower short-term barriers to initiating disclosure or early help-seeking activity when stigma-related concerns are prominent [11].

Empirical evidence supports the relevance of self-stigma to attitudes toward AI-mediated psychotherapy. Hoffman and Oppert [11] found that help-seeking self-stigma was associated with young adults' attitudes toward AI-delivered psychotherapy and their willingness to use a chatbot for therapy-like support. Their findings are useful because they indicate that self-stigma does not simply predict avoidance of all help [11]. Instead, it may shape which forms of help feel tolerable, acceptable, or "safe enough" to approach. This pattern is consistent with broader attitude research suggesting that chatbot support is often perceived as accessible and discreet, even where confidence in its therapeutic benefit is mixed [14]. In a survey of US university students, Rackoff and colleagues [14] reported that chatbot use for mental health support was uncommon, but chatbots were perceived as less constrained by practical barriers and stigma-related concerns than traditional services, while also being viewed as less beneficial overall. This combination is clinically interpretable: lower perceived exposure may increase openness to initial engagement, while uncertainty about efficacy limits reliance or sustained preference when other options are available [14].

Related constructs, particularly label avoidance, help explain why low-exposure tools can attract early engagement while leaving onward care uncertain. Label avoidance refers to reluctance to engage with mental health services because doing so is experienced as accepting a stigmatised identity and the social consequences attached to that label [14]. In other words, the barrier is not only anticipated judgement from others, but the meaning of being someone who "needs help", which can amplify avoidance and delay even in the presence of distress [14]. In this context, conversational tools can feel tolerable because they allow disclosure-like behaviour with lower immediate interpersonal exposure, without requiring the person to cross the social threshold of formal help-seeking [14].

Kosyluk and colleagues' screening-and-referral chatbot study provides a concrete example of this mechanism in action. In their US survey and prototype evaluation, label avoidance was the only significant predictor of chatbot use among participants currently experiencing distress, and most users who engaged with the chatbot completed screening and rated the system as acceptable [14]. This is clinically informative, but it should be interpreted precisely: the study supports the idea that a low-threshold chatbot can engage people who are motivated to avoid the "mental health patient" label, while leaving open the critical question of whether that engagement translates into later contact with human services [14]. Kosyluk, et al. [14] explicitly frame the work as feasibility and acceptability research and note generalisability limits related to sampling, reinforcing that pathway outcomes such as referral completion and time-to-care remain a priority for future studies [14]. Taken together, the evidence supports a bounded conclusion: self-stigma and label avoidance are not background attitudes, but mechanisms that shape which forms of support feel tolerable at the outset, and they may influence whether AI-mediated help-seeking functions as a bridge toward care or remains a private coping endpoint [23].

From chat to care: pathways, referral, and clinical plausibility

The central clinical question is not whether chatbots are efficacious in isolation, but whether they influence movement along pathways to care. In help-seeking research, movement from symptom recognition and appraisal to service contact is commonly disrupted at the point where an individual must translate distress into action, with internal barriers shaping whether intentions consolidate into help-seeking behaviour [3]. Barriers at this stage are frequently both practical and emotional, including service navigation demands, uncertainty about what level of support is indicated, concerns about stigma and disclosure, and the cognitive burden of decision-making under stress [3]. For this reason, conversational systems that incorporate screening, triage, and referral functions are of particular interest, because they are designed to reduce decision friction and navigation burden rather than provide generalised support alone [14].

Evidence for this pathway-oriented role is emerging. Kosyluk et al. [14] evaluated a mental health screening and referral chatbot and reported feasibility and acceptability, with the system tailoring resource recommendations to distress level and user preferences [14]. While feasibility studies do not establish effectiveness, they remain relevant to the pathways question because they test whether individuals will engage with a structured sequence that culminates in referral recommendations, rather than stopping at conversational support [14]. This reflects an empirically testable clinical hypothesis: reducing the cognitive and logistical demands of the next step may increase the likelihood of onward help-seeking among those who would otherwise delay or disengage at the transition from appraisal to action [3,14].

Clinical plausibility is strengthened by outcome evidence indicating that conversational agents can produce measurable reductions in distress in some contexts. In a systematic review and meta-analysis of AI-based conversational agents, Li, et al. [24] found evidence of symptom reductions in depression and distress and improvements in some wellbeing outcomes across experimental studies, with substantial heterogeneity that appeared related to intervention design and user experience [24]. These findings do not imply that any chatbot is therapeutically effective, but they do suggest that structured conversational interventions can influence symptoms under certain conditions, providing a rationale for examining their potential role as early supports or adjuncts while individuals navigate access to care [24]. More recent trial evidence extends this plausibility to fully generative systems. Heinz et al. [25] reported clinically meaningful symptom reductions from a generative AI therapy chatbot among adults with depression, anxiety, or high-risk eating disorder symptoms in a randomised design [25]. For pathways to care, the relevance is not only symptom change but the demonstration that conversational delivery can operationalise structured therapeutic content in a way that may be acceptable to users, leaving open the next question for the field: when, for whom, and under what conditions does engagement with conversational AI translate into onward contact with formal supports [14,25].

Practical translation: using AI as scaffolding, not a substitute

Given emerging evidence that people are already using generative AI tools for mental health related support in everyday contexts [1,6,7], a mechanism-consistent approach is to use conversational AI for tasks it plausibly supports at early stages of help-seeking: articulating symptoms, testing language for disclosure, and producing summaries that can be carried into human support when needed [4,5]. This stance is justified by two linked findings in the reviewed literature: first, disclosure comfort in chatbot interactions appears sensitive to perceived anonymity, fear of judgement, trust cues, and privacy expectations [4,5]; second, disclosure can be experienced as emotionally relieving and disclosure patterns can shift over time with repeated chatbot use, which makes a private "relief loop" a plausible failure mode even when users report benefit [4,13].

Where tools include structured screening and referral functions, feasibility evidence suggests that these flows can be acceptable and can reduce navigation friction, but linkage still depends on an explicit next step beyond the conversation [10]. Outcome evidence from trials and meta-analysis indicates that some conversational interventions can reduce distress under specific conditions, but this does not generalise to all tools or presentations, so persistent, worsening, or function-impairing symptoms remain an indication to escalate to human care [24,25]. One practical rule that follows from this evidence base is to convert chat into a shareable artefact (for example, an appointment summary or a drafted message) and complete one external action, rather than repeatedly re-narrating the same concern [10,13].

Gaps and further research

Current evidence supports the relevance of conversational AI to early help-seeking processes, yet several gaps limit confident translation into practice. First, much of the disclosure and acceptability literature relies on cross-sectional surveys, single-session experiments, or short follow-up windows, which constrains conclusions about longitudinal trajectories and real-world outcomes [12]. Longer prospective studies are needed to test whether AI-mediated disclosure predicts subsequent help-seeking behaviours, including appointment booking, referral uptake, adherence, and clinical engagement, rather than remaining primarily a private self-regulation tool. This is particularly important given evidence that disclosure patterns can shift over time within human–chatbot relationships [13].

Second, operational definitions vary. Studies differ in what they classify as "mental health use" of conversational AI, ranging from general advice-seeking to structured, therapy-like interaction, limiting comparability across samples and platforms [12]. Conceptual clarity and consistent measurement will be especially important when comparing large language model tools with purpose-built mental health chatbots, which differ in design intent and constraints [24,25]. Third, the pathway-to-care question remains under-tested. Feasibility work suggests screening-and-referral chatbots can be acceptable and can tailor recommendations, but this evidence does not yet establish whether these tools increase referral completion, reduce time-to-care, or improve clinical engagement [10].

Fourth, culture and context cannot be assumed uniform. Evidence indicates that emotional expression in chatbot interactions varies across cultural groupings, yet most studies do not directly test how cultural norms, language, and local service structures shape disclosure, trust, and subsequent care engagement [16]. Finally, clinical appropriateness and safety require finer-grained research in mental health terms, including clearer thresholds for when self-guided chatbot use is compatible with clinical need and when escalation pathways are required, particularly when distress is persistent, worsening, or impairing functioning.

Conclusion

This article reviewed emerging research on AI-mediated help-seeking in mental health by integrating evidence across three linked domains: self-disclosure in chatbot interactions, stigma-related barriers that shape uptake and delay, and the conditions under which conversational AI supports movement from private sense-making toward real-world care pathways. Rather than treating conversational AI as a single intervention, the review consolidated findings into a mechanism-led interpretation of where use is most likely to assist early help-seeking (for example, articulating distress, lowering perceived exposure, and reducing navigation burden) and where it can stall (for example, relief without action, trust-related limits on transfer, stigma bypass that maintains avoidance, and decision paralysis at the action threshold). The article's practical implication is that conversational AI is most defensible as early scaffolding when it produces transferable, action-oriented outputs that increase the likelihood of follow-through beyond the conversation, while recognising that current evidence does not support uniform benefit, does not resolve stigma as a construct, and does not guarantee linkage to care without an explicit next step.

References

1. Siddals S, Torous J, Coxon A. Experiences of generative AI chatbots for mental health: a qualitative study. Nat Ment Health. 2024;2:323-331. doi:10.1038/s44184-024-00097-4
2. Auf H, Svedberg P, Nygren J, Nair M, Lundgren LE. The use of AI in mental health services to support decision-making: scoping review. J Med Internet Res. 2025;27:e63548.
3. McLaren T, Peter LJ, Tomczyk S, Muehlan H, Schomerus G, Schmidt S. The Seeking Mental Health Care model: prediction of help-seeking for depressive symptoms by stigma and mental illness representations. BMC Public Health. 2023;23:69. doi:10.1186/s12889-022-14937-5
4. Croes EAJ, Antheunis ML, van der Lee ML, de Wit JBF. Digital confessions: the willingness to disclose intimate information to a chatbot and its impact on emotional well-being. Interact Comput. 2024;36(5):279-292. doi:10.1093/iwc/iwae016
5. Papneja H, Hasan H. Self-disclosure to conversational AI: a literature review, emergent framework, and directions for future research. Pers Ubiquitous Comput. 2024. doi:10.1007/s00779-024-01823-7
6. McBain RK, Bozick R, Diliberti M, et al. Use of generative AI for mental health advice among US adolescents and young adults. JAMA Netw Open. 2025;8(11):e2542281. doi:10.1001/jamanetworkopen.2025.42281
7. Ayre J, et al. Australians asking ChatGPT health questions: results from a June 2024 survey. Med J Aust. 2025. doi:10.5694/mja2.52598
8. Cross S, et al. Use of artificial intelligence in mental health care:

community and clinician perspectives. JMIR Ment Health. 2024. doi:10.2196/58462

9. Luo X, et al. Seeking emotional and mental health support from generative AI: a mixed-methods study of people who repeatedly use ChatGPT. JMIR Ment Health. Preprint. 2025.
10. Kosyluk K, Baeder T, Greene K, et al. Mental distress, label avoidance, and use of a mental health chatbot: results from a US survey. JMIR Form Res. 2024;8:e45959. doi:10.2196/45959
11. Hoffman BD, Oppert ML. Understanding young adults' attitudes towards using AI chatbots for psychotherapy: the role of self-stigma. Internet Interv. 2024;37:100722. doi:10.1016/j.invent.2024.100722
12. Papneja H, Yadav N. Self-disclosure to conversational AI: a literature review, emergent framework, and directions for future research. Pers Ubiquitous Comput. 2025;29:119-151. doi:10.1007/s00779-024-01823-7
13. Skjuve M, Følstad A, Brandtzæg PB. A longitudinal study of self-disclosure in human-chatbot relationships. Interact Comput. 2023;35(1):24-39. doi:10.1093/iwc/iwad022
14. Rackoff GN, Taliaferro LA, Borowsky IW. Attitudes and utilization of chatbot-delivered mental health support in a sample of US college students. JMIR Ment Health. 2025;12:e55854. doi:10.2196/55854
15. Xie Z, Xie Y. Would you share secrets with a chatbot? A longitudinal study on the relationship between virtual companionship and privacy disclosure. Aslib J Inf Manag. 2025.
16. Chin H, Yi J, Kim J, Hong JI. The potential of chatbots for emotional support and promoting mental well-being in different cultures. J Med Internet Res. 2023;25:e51712. doi:10.2196/51712
17. Sawma T, Hallit S, Younis CJ, Jaber J, El Khoury-Malhame M. Cultural stigma, psychological distress and help-seeking: moderating role of self-esteem and self-stigma. medRxiv. 2024.
18. Portt EL. Stigma associated with seeking help for psychological distress: how public and self-stigma, help-seeking attitudes and intent, self-compassion, and empathy relate. Doctoral dissertation. 2020.
19. Goetz Fidler CJ. Stigma resistance versus self-stigma in university students with mental health problems. Doctoral dissertation. 2023.
20. Pimenta SM, Hunter SC, Rasmussen S. The influence of stigma and trust in young people seeking support for their own or a friend's symptoms: the role of threat appraisals. Child Youth Care Forum. 2023;52:559-581. doi:10.1007/s10566-022-09698-6
21. Huang R, Lei B, Liu Y, Liu D, Zhang L. Self-stigma of seeking professional psychological help and its influencing factors among high-risk health care workers for depression and anxiety: a multicenter cross-sectional study. BMC Psychiatry. 2025;25:533. doi:10.1186/s12888-025-07019-4
22. Lacey L, Mishra N, Mukherjee P, et al. Can destigmatizing mental health increase willingness to seek help? Experimental evidence from Nepal. J Policy Anal Manage. 2025;44(1):97-124. doi:10.1002/pam.22643
23. Merikko J, Räsänen P. Exploring the role of AI guidance in internet-based acceptance and commitment therapy: links to human-computer trust and help-seeking barriers. In: International Conference on Human-Computer Interaction. Springer Nature Switzerland; 2025:221-239.
24. Li H, Zhang Y, Xie Z, et al. Systematic review and meta-analysis of AI-based conversational agents for promoting mental health and well-being. npj Digit Med. 2023;6:236. doi:10.1038/s41746-023-00979-5
25. Heinz MV, Mackin DM, Trudeau BM, et al. Randomized trial of a generative AI chatbot for mental health treatment. NEJM AI. 2025;2:AIoa2400802. doi:10.1056/AIoa2400802
26. Ennis E, O'Neill S, Mulvenna M, Bond RR. Chatbots supporting mental health and wellbeing of children and young people: applications, acceptability and usability. In: European Conference on Mental Health. 2023.
27. Mackenzie CS, Smith ML, Joubert L. Understanding help-seeking for mental health difficulties: a conceptual framework and future directions. Clin Psychol Rev. 2023;98:102231.