



EXECUTIVE INSIGHTS

When AI Enters the Therapy Room: From Workflow Automation to Hybrid Care Delivery

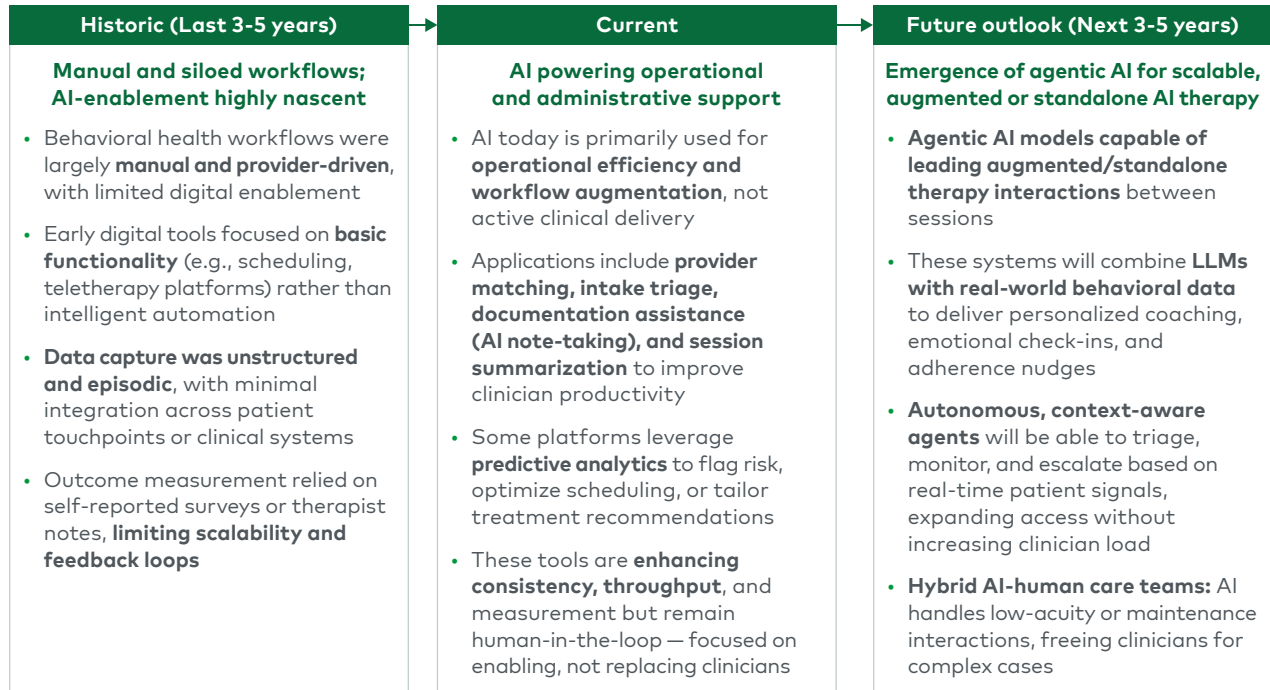
Executive summary

Behavioral health is emerging as one of the clearest proving grounds for applied artificial intelligence (AI) in healthcare. The need is obvious: Demand continues to rise; provider shortages remain persistent; and patients, clinicians, payors and employers all want more accessible, responsive and measurable care. At the same time, behavioral health is unusually sensitive to issues of trust, safety, privacy and human judgment. That combination makes the field both highly promising and uniquely demanding.

The practical implication is that the market is unlikely to move overnight from traditional care to fully autonomous AI-delivered therapy. The more credible path is a staged one. Today, AI is already creating value in provider matching, patient intake, documentation, risk flagging and between-session engagement. Over the next several years, the center of gravity is likely to shift toward hybrid models in which AI extends clinician reach, supports earlier intervention and improves continuity of care while licensed professionals remain central to diagnosis, escalation and management of higher-acuity needs.

The organizations most likely to shape the next phase of the market will not simply have the most advanced model. They will combine clinical grounding, thoughtful safety design, strong workflow integration, access to relevant real-world data, and the operational ability to fit AI into reimbursement, provider and care-delivery systems. In behavioral health, technical capability matters, but trust is what will determine scale (see Figure 1).

Figure 1
AI evolution in digital behavioral health



Note: AI=artificial intelligence; BH=behavioral health; LLMs=large language models
Source: L.E.K. research and analysis

AI is also beginning to reshape not just the care model of behavioral health but the business model. One reason is that behavioral health benefits are often distributed through overlapping sponsored channels rather than a single, exclusive payor relationship. In practice, the same individual may have access to behavioral health support through multiple benefit providers at once — for example, through an employer, an auto insurer or a life insurer — meaning the same underlying “life” can be monetized multiple times across the ecosystem.




Why behavioral health is such a strong use case for AI

Behavioral health sits at the intersection of high unmet need and strong digital suitability. Much of mental healthcare is delivered in outpatient settings, many interactions are conversational and longitudinal, and patients often value the privacy and convenience of virtual access. Those characteristics make the category more naturally compatible with technology-enabled care than most clinical domains.

The access problem is equally important. Demand for therapy, psychiatry and lower-acuity emotional support continues to outpace the available workforce. Wait times remain long in many markets, clinician availability is constrained and geography still shapes access in ways that are not clinically rational. AI will not solve the workforce shortage by itself, but it can help expand effective capacity. That is especially true when it reduces documentation burden, improves triage, keeps patients engaged between visits and supports lower-acuity needs that do not always require a full clinician encounter.

Economic pressure reinforces the same conclusion. Payors want earlier intervention and lower downstream medical spend. Employers want benefits that improve access, reduce disruption and produce usable outcome indicators. Clinicians want tools that reduce administrative friction without compromising therapeutic judgment. Patients want care that is easier to start, easier to continue and easier to fit into daily life. In a category where all stakeholders are asking for some version of greater access and efficiency, AI has clear structural relevance across the system (see Figure 2).

Figure 2
Digital behavioral health stakeholders and their needs

Key stakeholder	Key needs/pain points	Drivers of digital BH solution adoption
 Customers (i.e., payors and employers)	<ul style="list-style-type: none"> • Limited visibility into care quality and outcomes and difficult to measure ROI • Access gaps driven by provider shortages and condition-specific coverage gaps (e.g., SUD, ABA) • Fragmented coordination between payors, PCPs, and behavioral providers 	<ul style="list-style-type: none"> • Broader network access and faster member connection to care • Measurable outcomes and clear reporting of clinical improvement/ROI • Seamless integration and potential for downstream cost savings
 Providers	<ul style="list-style-type: none"> • Inconsistent patient flow and variable reimbursement stability • High admin burden from scheduling, charting, billing, and credentialing • Limited data insights and tools to track outcomes or manage panels 	<ul style="list-style-type: none"> • Expanded patient reach and stable, predictable caseload • Streamlined workflows and reduced admin time via digital tools • AI-enabled clinical support (e.g., note-taking, patient insights, scheduling)
 Patients	<ul style="list-style-type: none"> • High cost and insurance friction create barriers to sustained care • Long wait times and limited access to in-network providers • Difficulty finding the "right fit" therapist; privacy and stigma concerns 	<ul style="list-style-type: none"> • Covered care on a recognized platform • 24/7 virtual access and flexible scheduling across modalities • Personalized matching, anonymity, and convenience enhance engagement

Note: BH=behavioral health; ROI=return on investment; SUD=substance use disorder; ABA=applied behavior analysis; PCPs=primary care physicians; AI=artificial intelligence
 Source: L.E.K. research and analysis

Where AI is creating value today

Despite the excitement around generative AI, the current state of adoption in behavioral health is still more operational than clinical. The most established use cases sit behind the scenes rather than at the center of clinical care. AI is being used to support intake and triage, improve provider matching, generate notes and session summaries, flag signs of risk or dropout, optimize scheduling, and personalize reminders or self-guided prompts. In other words, the first wave of value creation has focused on making the existing care model work better. Increasingly, these tools are also commercial infrastructure; by improving digital triage, self-service engagement and referral conversion, they help vendors capture value at more than one point in the member journey rather than relying on a single reimbursed encounter.

That first wave will be more foundational than it may seem. Better matching can reduce friction at the point of entry. Better documentation can give time back to clinicians. Better monitoring can identify risk earlier and make follow-up more targeted. Better data capture can improve outcomes visibility for organizations that need to understand whether a mental health intervention is actually changing behavior, adherence or utilization. None of these use cases depend on replacing the therapist or psychiatrist. They improve performance by tightening the system around the clinician and the patient.

The market is beginning to move beyond those enabling functions, but cautiously. During 2025 and early 2026, several public examples helped move the conversation from abstract promise toward more concrete evidence and governance. Spring Health released VERA-MH, an open-source evaluation framework for mental health chatbots. Lyra began piloting a clinically designed AI guide with risk-flagging and escalation pathways. Dartmouth College researchers published early clinical-trial results for Therabot. OpenAI convened an Expert Council on Well-Being and AI alongside related clinical advisory work. Those developments matter because they suggest the field is shifting from experimentation toward evaluation, governance and evidence generation. Still, the core message remains the same: Current momentum is real, but broad acceptance for AI-led mental healthcare has not yet arrived.

Why hybrid models are likely to win first

The next chapter is likely to be defined by augmentation rather than immediate autonomy. Hybrid models are more acceptable because they preserve a role for human judgment where the stakes are highest while still unlocking some of AI's strongest advantages: always-on availability, lower marginal cost, more frequent touchpoints and the ability to synthesize large amounts of behavioral data over time.

In practical terms, that means AI is more likely to succeed first as a copilot rather than as a substitute. It can reinforce skills between sessions, prompt journaling or check-ins, help identify relapse risk, summarize patterns for clinicians, support waitlist management and provide lower-acuity guidance in bounded contexts. Meanwhile, clinicians remain responsible for assessment, diagnosis, high-risk situations, medication decisions and changes in treatment strategy. This model aligns better with current trust dynamics and current reimbursement realities, which still tend to favor clinician-led care even when technology is deeply embedded in the experience.

Selective autonomy may still emerge, particularly for low-acuity support, maintenance use cases or early engagement before a patient enters formal treatment. In those settings, the appeal is obvious: AI can be available at any hour, engage patients who might otherwise defer care, and create a bridge between symptom recognition and more structured support. But autonomous use cases will scale only where safety expectations are explicit, escalation pathways are reliable, and the model operates inside clearly defined clinical and ethical boundaries (see Figure 3).

This distinction matters strategically. The debate should not be framed as “AI therapist versus human therapist.” The more useful questions are: Where can AI safely add value across the behavioral health journey? What level of autonomy can each use case support? What evidence is required before stakeholders will trust it?

Figure 3
Augmented vs. stand-alone strategy

	AI for care augmentation	AI for standalone therapy
Description	<ul style="list-style-type: none"> • Clinician-supervised AI that supports patients between sessions (e.g., progress monitoring, engagement) while clinicians retain decision-making (e.g., diagnosis, treatment plan) 	<ul style="list-style-type: none"> • AI manages the full care journey with varying degrees of clinician oversight* — from diagnosis to follow-up — with risk-flagging and escalation protocols in place
Benefits	<ul style="list-style-type: none"> • Perceived safety: Clinician oversight provides trust and accountability/monitoring • Provider capacity scalability: Extends clinician capacity (larger caseloads) and reach 	<ul style="list-style-type: none"> • Scalability: Reaches more patients clinician-led models (which are limited by provider count) • Cost: Far lower cost than human-led or hybrid modalities
Customer receptivity	<ul style="list-style-type: none"> • Payors & Employers: Both segments are cautious but more open to AI with clinician guardrails and clinician involvement • Practices (BH and non-BH): Receptive to augmenting AI that extends capacity without loss of control or oversight 	<ul style="list-style-type: none"> • DTC: Most receptive, particularly younger patients (teens, young adults) • Payors: Currently apprehensive, given safety, intervention efficacy, and risk concerns, without a clinician in the loop; less agile than other groups • Employers: Similarly apprehensive as payors, but likely more agile • BH practices: Wary of job displacement (cannibalization) and safety risks • Non-BH practices: More likely to adopt to solutions to supplement their care (e.g., PCP, specialty), but will require safety and efficacy validation

Augmented vs. stand-alone strategy

What safe scale will require

If the market moves from workflow automation toward more clinically active AI, the gating question will be about safety at scale. In behavioral health, it is not enough for a model to sound empathetic or to generate plausible therapeutic language. It must perform reliably in moments of ambiguity, emotional volatility and potential crisis. That requires a stronger operating model than many general-purpose AI deployments have needed. Unlike many AI applications, failure in behavioral health carries asymmetric risk. Errors are far from just inconvenient; they are potentially harmful, thereby raising the bar for deployment.

A practical way to think about this involves the six conditions for success.

- 1. Trust:** Users, clinicians and institutional buyers need confidence that the system is clinically grounded, transparent and used in ways that are easy to understand.
- 2. Data strength:** The model must be informed by relevant, high-quality behavioral health data and evaluated against real-world outcomes, not just general language benchmarks.
- 3. Accessibility:** The solution must be available in the places where people actually seek support, including payor-backed and employer-sponsored channels, and within a cost structure that makes it broadly accessible.
- 4. Personalization:** Useful mental health support is inherently contextual, so the experience must adapt to the individual rather than delivering generic advice.
- 5. Safety:** Crisis recognition, escalation, boundary management, bias monitoring and auditability are not optional design features. They are core infrastructure.
- 6. Ecosystem integration:** AI must fit within provider workflows, documentation expectations, care pathways and reporting requirements rather than sitting outside the system as a disconnected tool.

Taken together, these conditions suggest that success in behavioral health AI will not come from model sophistication alone. It will come from combining product capability with clinical operations, governance and distribution. This is one reason many early market signals have focused so heavily on safety councils, evaluation frameworks, guardrails and human oversight instead of on raw model performance (see Figure 4).

Figure 4

Characteristics of a successful AI model in behavioral health

**Engenders patient, provider, customer, and regulator trust through clinical grounding**

- Trust comes from clinical grounding, transparency, explainability, and ethical guardrails — especially in sensitive mental health use cases
- Patients and clinicians must perceive AI as an augmenting partner (not a replacement) in the care ecosystem, reinforcing confidences in outcomes and privacy

**Trained on extensive “real” clinical data (not synthetic or generic foundational models)**

- AI must be trained on large volumes of real, clinically validated therapy and engagement data, not synthetic text; access to longitudinal, multi-modal data (e.g., transcripts, clinical notes, outcomes) enables accurate, context-aware guidance and risk identification

**Widely accessible (including both availability and cost effectiveness)**

- AI should be embedded within existing BH platforms and payor networks, ensuring broad patient reach across coverage types and reimbursement structures
- Cost-efficient scalability is critical to support high-frequency, low-acuity interactions

**Highly personalized and easy-to-use for target users (i.e., patients) with multi-modal capabilities**

- Models must adapt to individual patient needs, context, and emotional tone, mirroring therapeutic best practices
- Further, simple, intuitive interfaces improve adherence and sustain engagement, particularly for anxious or low-digital-literacy users

**Built with safety guardrails to avoid adverse events for target users**

- Built-in risk detection and escalation protocols are essential to identify crises (e.g., suicidality, self-harm) and route to clinicians in real time
- Continuous supervision and clinician-in-the-loop design ensure ethical, compliant delivery and mitigate liability risk

**Owned by platform vertically integrated across the solution ecosystem (i.e., BH landscape)**

- AI must be anchored in a platform that owns the full BH workflow — from payor contracting and clinician network to data collection and outcomes tracking
- Integration across stakeholders (payors, employers, clinicians, patients) creates feedback loops that improve accuracy, safety, and ROI over time

Note: AI=artificial intelligence; BH=behavioral health
Source: L.E.K. research and analysis

How the patient journey will change

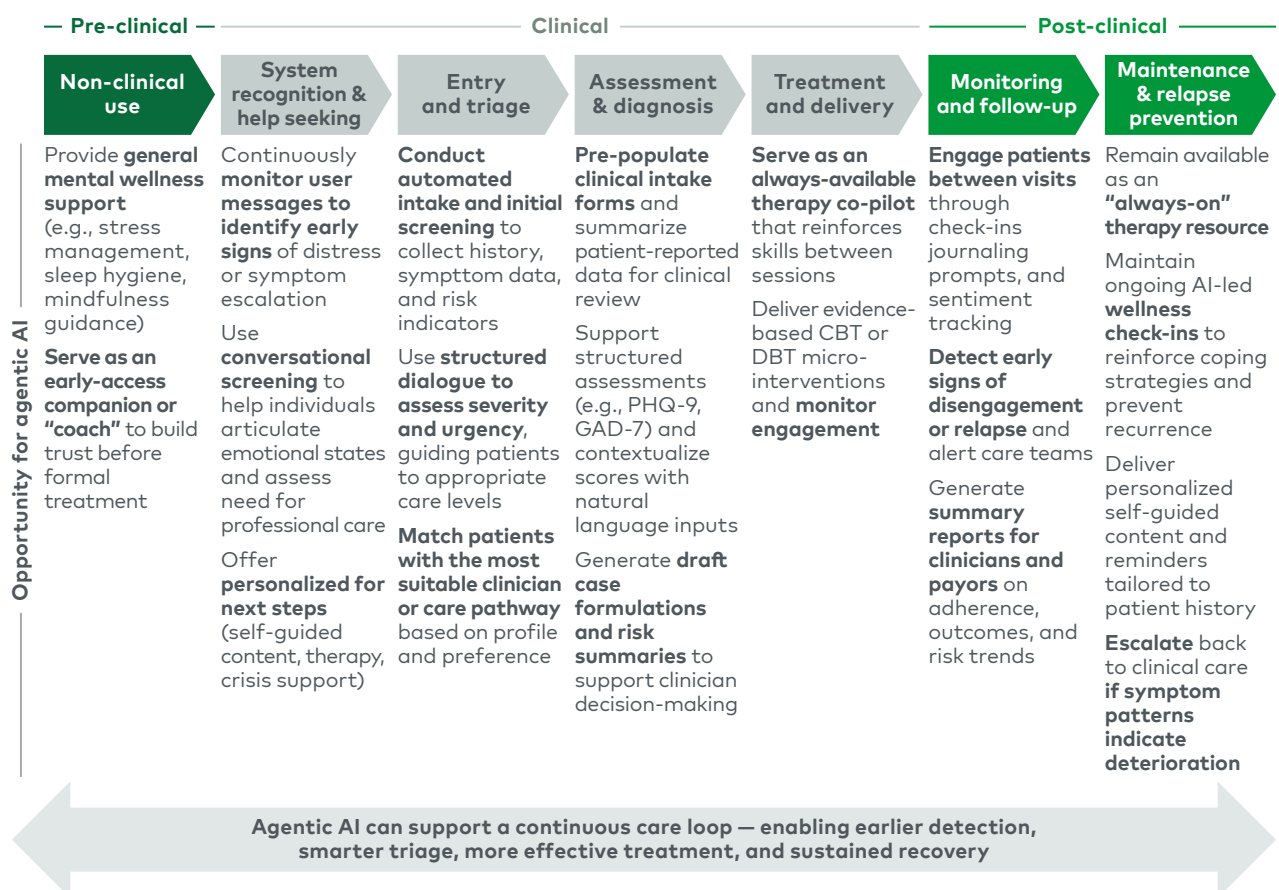
The most useful way to understand AI's future role is not by model type but by patient journey. The deeper shift is from episodic care to continuous support. Traditional therapy is structured around discrete sessions. AI introduces a persistent layer of monitoring, engagement and intervention between those moments. At the front end, AI can help people articulate what they are feeling, complete screening and navigate toward the right level of care. AI may also reshape how patients enter care, serving as a low-friction, always-available first touchpoint for individuals who might otherwise delay or avoid seeking help. During entry and triage, AI can collect history, structure symptoms and improve routing. During treatment, it can reinforce skills, prompt reflection, surface changes in mood or engagement, and support continuity between scheduled encounters. After more acute treatment, it can play a role in relapse prevention, maintenance and re-entry into care when signs of deterioration reappear.

These are meaningful changes because behavioral health has always struggled with fragmentation between moments of active care. Traditional models often have limited visibility into what happens between appointments, limited ability to intervene early when symptoms worsen and limited capacity to maintain lower-intensity support over time. AI can help create a more continuous support layer around the formal care relationship. That continuity may prove to be one of the most valuable aspects of the technology, particularly for patients whose needs fluctuate and for clinicians managing larger caseloads.

Importantly, not every stage of the journey carries the same clinical risk, which is why adoption will almost certainly be uneven. Lower-acuity screening, behavioral coaching, journaling prompts and relapse monitoring are more likely to move first. High-risk assessment, crisis response, diagnosis and medication management will continue to require a much higher threshold of oversight and evidence. The market is unlikely to converge on one universal model of AI-enabled behavioral healthcare. It is more likely to segment by acuity, use case and care setting (see Figure 5).

Figure 5

Specific use cases for agentic AI in the behavioral health patient journey



Note: AI=artificial intelligence; BH=behavioral health; PHQ-9=Patient Health Questionnaire-9; GAD-7=Generalized Anxiety Disorder 7-item scale; CBT=cognitive behavioral therapy; DBT=dialectical behavior therapy
 Source: L.E.K. research and analysis

How AI is changing the business model, not just the care model

One of the more important developments in behavioral health is that AI is not only expanding what can be delivered digitally; it is also broadening how vendors get paid. Historically, many behavioral health solutions relied on a relatively simple model: direct-to-consumer subscriptions, employer-paid access fees or reimbursed clinical encounters. That model is becoming more layered. As AI tools are embedded earlier in the patient journey and across more channels, revenue increasingly can be captured at multiple points — at the population level through employer or health plan contracts, at the engagement layer through digital navigation and self-guided support, and at the clinical tier when a member steps up into coaching, therapy, psychiatry or other covered services. These models are not always mutually exclusive at the life level; in practice, the same individual may sit inside multiple sponsored benefit channels at once, so the same covered life may be monetized more than once across the ecosystem.

This is creating a model that looks increasingly familiar to anyone who has followed employee assistance program (EAP) evolution: A broad population may be covered through an embedded or access-based contract, but only a subset engages deeply enough to generate additional utilization, referral or care-management revenue. AI can make that structure more powerful by serving as the always-on front door — identifying need earlier, improving engagement, routing people to the right level of care and keeping members active between visits. In that sense, AI is not only a clinical tool; it is also a conversion, retention and care-navigation layer. The same covered life may therefore be monetized sequentially across multiple offerings, especially when the vendor sits inside the carrier, employer or provider workflow. Rather than a standardized industry rule, this is an inference from how many current offerings are being structured.

The market is already moving in this direction. Spring Health pairs an employer-facing EAP+ model with broader platform services and has also expanded its behavioral health platform to more health plans. Lyra markets a single platform across employers and health plans, combining guided self-care, provider matching, coaching, therapy, medication support and AI-enabled tools. Headspace continues to position itself with health plans through a mix of content, digital support and provider-backed care. Wysa, meanwhile, spans employers, insurers, health systems and primary-care-linked pathways, using AI self-help and navigation to engage users and then route them onward to other services when appropriate. The implication is that the winning business models may be those that combine low-cost, population-level reach with the ability to capture higher-value clinical experiences, navigation assistance and workflow revenue as needs intensify.

For buyers, that makes vendor evaluation more complex. It is no longer enough to ask whether an AI tool “works.” The better question is whether the vendor’s commercial model aligns incentives across access, engagement, escalation and outcomes. Solutions that monetize only at the point of high-acuity clinical use may struggle to justify broad deployment. By contrast, models that combine access fees, embedded distribution, and step-up care revenue may have stronger economics and more room to invest in safety, evidence generation and workflow integration over time.

What could slow adoption

The barriers to adoption are real and should be treated as structural rather than temporary. Clinical risk is the most obvious one. Behavioral healthcare includes situations involving suicidality, other self-harm, trauma, psychosis and rapid decompensation. A system that performs well most of the time but fails badly in edge cases is not ready for broad deployment. Safety, therefore, is not just a product issue; it is a market access issue.

The evidence base is another limiter. Early signals are promising, and formal research activity is increasing, but many stakeholders will still want more proof before they treat AI-enabled behavioral support as clinically interchangeable with human care. Dartmouth’s early Therabot trial is an important indication that more rigorous evidence is starting to emerge, but one encouraging study certainly does not indicate a broad clinical consensus. Payors, in particular, are likely to remain cautious until evidence, practice guidance and payment pathways mature further. Employers may move somewhat faster in lower-risk use cases, especially as more guided-support pilots come to market, but they too will need confidence that innovation is not outrunning governance.

Provider acceptance will also shape the curve. Clinicians are broadly receptive to tools that reduce burden and improve continuity, but many remain skeptical of tools that threaten therapeutic quality, professional responsibility or the therapeutic alliance itself. Their skepticism is not an obstacle to be brushed aside. It is a useful design constraint. The strongest products will be the ones that make clinicians more effective, not the ones that assume trust can be bypassed.

Finally, the policy environment is beginning to move from abstract concern to real governance. As formal frameworks, state-level reviews and best-practice guidance develop, the field will become easier to evaluate but harder to bluff. Utah’s 2026 regulatory experience is one example of how quickly the discussion can shift from general AI enthusiasm to concrete inquiries around disclosures, escalation and consumer protection in mental health settings. That is a healthy shift. In a domain this sensitive, discipline is likely to be a precondition for adoption rather than a brake on it. Beyond clinical risk and evidence, reimbursement remains a gating mechanism. Until

AI-enabled interventions are recognized within formal payment models, adoption will remain constrained regardless of technical progress.

What to watch over the next three to five years

The most likely near-term outcome is that augmentation will become the norm before autonomy becomes mainstream. AI will become more common in intake, engagement, measurement, documentation and lower-intensity support. At the same time, a smaller set of bounded autonomous use cases will continue to be tested, especially in lower-acuity settings or in direct-to-consumer channels where the need for always-available support is especially visible.

Several signals will determine whether the market moves faster. One is evidence: More clinical trials, stronger real-world outcomes, and clearer differentiation between low-risk wellness tools and clinical-grade interventions are required. Another is regulation: This means practical standards for disclosure, escalation, documentation and accountability — not necessarily heavy-handed regulation. A third is reimbursement: Once payment models better recognize AI-supported or AI-mediated care, adoption could accelerate meaningfully. A fourth is institutional trust: As more buyers see rigorous guardrails, transparent evaluation and credible oversight, it will become easier to move from pilot to scaled deployment.

For the broader market, the implication is clear: AI will matter in behavioral health, but the issue is not simply whether it works. The real questions ask where it can be trusted, how it will be governed and which organizations can combine product capability with clinical credibility. The market is now moving out of the broad curiosity phase and into the phase of selective proof. That is exactly where serious categories begin to take shape.

A related implication is that the most attractive behavioral health models may be those with multiple monetization layers around the same covered population. In practical terms, that can mean pairing carrier-embedded or employer-paid access revenue with AI-enabled navigation, step-up referral, claims-based clinical care, and in some cases provider-facing workflow tools or analytics. As in EAP, breadth of covered access matters — but the real strategic advantage may come from owning the member journey across multiple points of engagement and escalation.

Conclusion

Behavioral health is one of the clearest real-world tests for AI: Can it improve care in a way that is both scalable and responsible? The need is urgent, the economics are supportive and the digital fit is stronger than in many other areas of healthcare. But it is also a category in which harm, if it occurs, will be highly visible and deeply consequential.

That is why the future is unlikely to belong to the boldest claims. More likely it will belong to the most disciplined models – those that improve access without overreaching, extend clinician capacity without eroding judgment, and create more continuous support without losing sight of safety, evidence and trust. Over the next several years, the center of the market will likely move toward AI-enabled hybrid care. The organizations that succeed will be the ones that treat trust not as a communications issue but as part of the product itself. Over time, the sector's leaders may be distinguished not only by clinical outcomes and safety, but also by their ability to build resilient, multilayered business models around access, engagement, escalation and longitudinal member value.

For more information, please [contact us](#).

About the Authors



Matt Sabbatino

Matt Sabbatino is a Managing Director and Partner in L.E.K. Consulting's New York office and a leader in the firm's Healthcare Services practice. Matt advises clients on growth strategy, M&A, performance improvement, organizational transformation and value creation across the healthcare ecosystem. He has expertise in behavioral health, health plan strategy and operations, payer services, and the broader employer benefits space. His clientele includes health plans, healthcare services organizations, provider platforms, healthcare IT and digital health companies, and PE investors.



Harsha Madannavar

Harsha Madannavar is a Partner in L.E.K. Consulting's Technology Infrastructure practice globally (software, data, services and silicon systems), advising US and global corporates, private equity and hedge funds on a range of shareholder value issues including growth strategy, business model transformation, technology disruptions, product development, pricing and M&A. Harsha is a member of the firm's global Board of Directors and also serves on the Board's Audit and Risk committees.



Brendan Mitchell

Brendan Mitchell is a Managing Director and Partner in L.E.K. Consulting's Healthcare Services practice and is based in the firm's San Francisco office. Brendan focuses on growth strategy, M&A and operations-related topics and has used his experience to serve leading national and regional payers, risk-bearing provider organizations, healthcare IT firms and private equity investors. His core areas of expertise include care model design, clinical operations and healthcare payments.



Vidhu Venkatesh

Vidhu Venkatesh is a Senior Manager in L.E.K. Consulting's Los Angeles office within the Healthcare services practice. Vidhu supports clients across the HCS ecosystem, spanning provider organizations, payers and healthcare IT providers on strategic planning, competitive positioning, growth strategy and transaction support. He specializes in behavioral health and has advised clients on market entry, network adequacy, value-based care models, AI strategy and the evolving reimbursement landscape across mental health and substance use disorder services.

About L.E.K. Consulting

L.E.K. Consulting is a premier global strategy consulting firm that partners with clients to solve their most critical business challenges and deliver high-impact outcomes. We bring deep industry expertise and rigorous, evidence-based insight to global corporations, growth companies, and private equity investors, helping them turn strategy into action. Founded more than 40 years ago, L.E.K. now operates across 27 offices worldwide. Learn more at lek.com.

L.E.K. Consulting is a registered trademark of L.E.K. Consulting LLC. All other products and brands mentioned in this document are properties of their respective owners. © 2026 L.E.K. Consulting LLC