

Webinar Title: The Adaptive Mind: How AI Changes the Way Clinicians Think

Webinar Date and Time: Friday, January 30th, 2026, at 1PM EST

Jodi: Hello, and welcome to today's webinar, The Adaptive Mind: How AI Changes the Way Clinicians Think. This topical webinar is paid for by Upheal, a secure AI-powered progress notes and analytics platform designed by clinicians to save mental health providers time. Some important points before we get started. Ted Faneuff, the presenter for today's webinar, has a working relationship with Upheal, and as such, he has both a financial and intellectual interest in the webinar content.

The purpose of the presentation is to provide a balanced view on how AI changes the way clinicians think. APA does not endorse any products or services, and the content was created by Upheal and does not reflect the views of APA or its editorial staff. Email advertising@apa.org to learn more. Next, this program does not offer CE. However, we will email everyone watching live today a certificate of attendance. Certificates will only be issued to those who watch for a minimum of 45 minutes.

A recording of this presentation will be emailed to everyone in two weeks' time, and the recording will include presentation slides. During our time together, you will be on mute, and you can communicate with us using the Q&A box located on your webinar screen. The Q&A is enabled for attendees to interact with submitted questions.

Attendees can view all submitted questions and upvote those they find important. This helps the presenter identify questions that many attendees want answered. A link to the presentation slides will be posted in the chat box of your webinar screen, and if you miss them, don't worry, we'll email them to you with the recording. Now, let me introduce you to Ted Faneuff. Ted?

Ted Faneuff: Good afternoon, everyone. Like Jodi had said, I'm Ted Faneuff. It's really a privilege to be here today for this webinar. Just to give you a little bit about my background, I'm a licensed clinical social worker, and I'm a practitioner. I currently lead Clinical Operations at Upheal, which is a platform harnessing AI to help therapists in their workflows. I've been talking about the concept of integrating AI into clinical workflows ethically and responsibly for the last two years.

I think it's truly an exciting time that we're living in as mental health providers. Contrary to even a few years ago, we're entering into an era where technology really has the capability to augment the work that we do in meaningful ways. AI isn't just a separate tool that we're occasionally opening now and asking our everyday questions; it's really becoming an ambient presence in our workflows and our daily lives. We see it summarizing our sessions, flagging risks, and really assisting in assessments, even in some cases.

While the industry has so far been focusing on how to use these tools, today we're going to look beneath the surface, and we're going to be discussing a concept I'm referring to or calling the adaptive mind. Our goal is to understand how working alongside AI or technology in general is fundamentally reshaping our cognitive processes, our memory, our curiosity, and our clinical intuition.

As we move through the next hour, I really want to challenge you to think of your mind as an active participant in this human AI system, and I want to explore how we're adapting intentionally and otherwise to the logic of the software that we use every day. Some of what you're going to hear today is going to feel repetitive, and I want to tell you upfront that's by design. A mentor of mine likes to say that leaders are repeaters because then things stick, things absorb.

Because the topic we're discussing, how AI reshapes clinical thinking, it doesn't show up as one dramatic aha moment, I think it actually shows up as a set of small, predictable shifts or ideas that repeat across contexts. As we move through the talk, I'm going to keep returning to some similar anchors, not because I'm stuck or I'm glitching or anything like that, but because this is exactly how cognitive drift works. It's gradual, it's quiet, it's easiest to miss when each individual moment feels reasonable.

Here's the core points that we're going to be circling today. First, AI changes more than workflow. It actually changes how we notice, prioritize, and close the loop on meaning. Second, we're going to be looking at-- we're practicing in constrained environments. Time pressure, documentation load, productivity demands, everything smooth and fast becomes cognitively persuasive. Third, the main risk isn't obvious error here. It's subtle drift. It's less checking, fewer alternatives, quicker closure because the story feels fluent.

Then fourth, the drift isn't moral failure. This is predictable adaptation under constraint. Then finally, this is predictable. It's manageable. Through small scaffolding, we can keep monitoring curiosity and clinical judgment online. If you hear me repeat these things, this is the intervention here. We're training attention to notice the quiet patterns and the things that are important.

Next slide, please. When we talk about AI and clinical work, it's tempting to keep the conversation at the level of permission. Is this allowed? Is it compliant? I think that those questions are certainly important, but they're not the most psychologically interesting question anymore. I think the more clinically consequential question that we're coming up against is, what does working with this kind of system do to the way that we think as clinicians?

The most important effects of AI usually won't announce themselves as ethical breaches or obvious errors. They're going to show up as these quiet shifts that occur. How we're encoding a session in our memory, how quickly we experience closure around a formulation, how long we tolerate uncertainty before we reach for an explanation, and what we treat as salient enough to revisit later. In other words, AI may not change our values, but it can change our epistemic habits.

I think that that's super important in psychotherapy because so much of our effectiveness rests on our capacities that are easy to compress when we don't notice. That sustained attention, that curiosity, that discomfort tolerance, the hypothesis testing, and the ability to hold those competing interpretations without prematurely collapsing into a single narrative. I want to name a core premise upfront, that human cognition is adaptive. We really are built to economize the effort in the environments that offer structure and cues.

I think the question isn't, will clinicians adapt to AI? We are, and we will. I think the question becomes whether that adaptation is intentional, which is grounded in our reflective choices, or if it's passive, where the workflow quietly trains our minds by rewarding speed, neatness, and completion. Today is really about making those shifts visible so we can design and use these tools in ways that protect what clinical expertise really depends on, which is our reflective judgment, not just efficient output.

Next slide. Most AI conversations right now have gone on in mental health really start with a tool layer. What does it automate? What does it summarize, score, or flag? Again, those are important because the practical realities of the workload and documentation are really very real in how we exist right now. For mental health professionals, there's really a deeper layer that's easy to miss if we stay only at the feature level.

It really is, when a system consistently highlights certain things, it doesn't just help us. It shapes what we notice. It shapes what we store, what we return to. In cognitive terms, tools don't only reduce effort, they're reorganizing that attention and the salience, the things that feel clinically important in the moment, what feels complete after the fact. That completion piece is subtle, but it's huge in psychotherapy because a good session often ends with something still open.

When AI produces a coherent narrative quickly, especially if it's a clean summary, it can create a premature sense of closure. Again, it might not be that it's wrong, but it's tidy. Tidiness is cognitively seductive. It feels good at the end of a session or day to just be done. We're going to keep returning to this question today, but what does working alongside AI do to our minds? Not in a dramatic, oh, it's taking over, prepare for the robo-rebellion of AI therapists, but in quiet ways.

The ways that you might see it are what we rehearse, what we forget, where do we stop double-checking, what do we stop lingering on? I think an important thing is to understand how this might land differently depending on experience level, because the tool is really shaping different parts of the craft as we're working with it.

For trainees, that might look like AI acting like training wheels. It can stabilize the early phase of practice by making the hour, maybe feel more organized and legible. The risk isn't that they become dependent on the tool. It's that their attention gets trained on whatever the system actually reliably names, and the unlabeled parts of the hour get the fewer reps to increase their judgment.

For experienced clinicians, the question isn't whether they can think without it. We certainly can. The question is whether the tool quietly redefines what counts as enough to us, enough to move on. When something feels synthesized already, the threshold for pausing or doing that work of recheck or staying with the unresolved can creep upward over time.

I want to also point out that this isn't really an argument against AI. That would be counter to the work that I do. It's an argument more for metacognitive awareness, noticing what the tool rewards in our attention, and intentionally protecting the parts of clinical reasoning that don't naturally survive optimization.

Next slide. If you're currently using AI in any part of your clinical work, whether it's documentation, summaries, et cetera, there's a good chance that this slide might feel familiar. What a lot of clinicians describe is the same initial experience. They notice that they're faster, they spend less time on notes, and they might feel more organized. I want to say really plainly that those are really legitimate gains we have to take into consideration.

We really shouldn't pretend that administrative burden is trivial because burnout and moral injury aren't abstract in this field. They really are here. Reducing that friction matters. Mental health professionals are really trained to look for what else changes when something gets easier. When a tool reduces effort, it doesn't only reduce workload, it changes where the work happens, and sometimes whether that happens at all.

Alongside speed and organization, many clinicians report something quieter, not that reflection vanishes, but that it moves. It becomes more after the fact, and it's more shaped by what the system captured and more likely to follow the structure that the tool provides. The clinical mechanism that I want you to hold onto is when information is externalized, summarized, categorized, made more retrievable, our minds often do less internal replay and less internal checking.

For us in psychotherapy, that replay is not just memory. It's how we digest the session. The things that we ask might normally be, what did I miss emotionally? What hypothesis am I leaning towards too quickly? What felt clinically hot that I didn't follow? What do I need to bring to consultation? If the summary becomes the most convenient representation of the hour, it slowly can become an authoritative representation. Again, I want to say that sometimes it's very accurate. It's not because it's inaccurate, but because it's cognitively easy. What's easy to retrieve often becomes what feels true enough to move on.

Next slide. Before we talk more about risk, I want to ground this in something cognitive psychology has been saying for a long time. I'm not a cognitive psychologist. I just did a lot of research on this, so I won't take that away from any of the cognitive psychologists in the group. Cognitive offloading is really not a bug in our human thinking. It's actually a feature of it. We really have always used environment to support our minds.

If you think about the work before AI or telehealth, we used notes, calendars, checklists, supervision notes. Even the way we structured the session itself helped us stay grounded in what we were doing. Offloading really reduces cognitive load, and in many settings, it does improve performance. If the goal is simply remember more or keep track of more, external supports really can be genuinely helpful.

Here's the hinge. Offloading changes where responsibility lives. When you offload something, you're not just storing information somewhere else; you're deciding, often implicitly, which part of the task your mind no longer needs to hold with the same intensity. In psychotherapy or the work that we do, the craft isn't knowing what happened. It's the continually sensing the pattern beneath it, including what's said, what's felt, what doesn't land, where your own confidence might be getting ahead of your work, et cetera.

AI obviously will increase the offloading opportunity dramatically because it doesn't just hold notes like a notepad. It processes. It summarizes. It selects what it counts as important, and it hands you back a version of the hour that already has an interpretive frame, and that's important. Again, we're not asking this question. It's not whether offloading is good or bad. We know that it's a natural thing in humanity.

I think more so, it's, do we still know what we're doing internally and what we've quietly handed to a system? A simple way to say that is offloading is adaptive when it frees the cognition that results in for deeper work or presence, but it becomes dangerous when it replaces that deeper work, and we don't even notice the swap. This slide, really, represents a transition point for the rest of this talk. We're going to map out what tends to get offloaded and how that changes clinician monitoring, and what design or use choices keep judgment active instead of outsourced.

Next slide. When we talk about offloading-- by the way, I apologize, I'm probably going to be going quick because I've got an hour, so I apologize if it's fast. When we talk about offloading, it's really easy to think that we're only talking about recall, which is what we remember from the session, but the more important shift cognitively is how often you monitor your work after the session.

The offloading literature is really clear on this. When an external support is dependable, people don't just remember differently, they allocate less effort to rehearsal and checking because the monitoring itself is costly. The environment is now doing part of the job. In prospective memory studies, once intentions are placed into reliable external cues, people can depend more on perceptual triggering than active internal monitoring, again, not because it's careless but because it's efficient.

If we translate that into psychotherapy practice, if a system gives me a clean structure, like a summary or a set of themes, the risk isn't that I'm unethical or unintelligent. It's actually that the first layer of my reflective scanning may arrive from the outside, rather than emerging from my own uncertainty and embodied sense of the hour that occurred.

Again, some of the questions, instead of my mind spontaneously doing the work of questioning, like what stood out that I can't quite name yet, or what might I be missing, I may start with what the system surfaced, and then I'm reflecting, but I'm reflecting within a frame that I might not have generated.

Here's the tight bridge to automation research. In clinical decision support contexts, automation bias is explicitly described as overaccepting computer output as a heuristic replacement for vigilant information seeking and processing. Those effects show up as both commission and emission patterns when our internal monitoring drops. It's really attentional reallocation under the conditions of perceived reliability, the time pressure that we work under, the workload, and a trust calibration over time.

In our work, the question really is what happens to clinical judgment when the starting point of judgment shifts from internally generated monitoring to externally provided cues? I think importantly, I want to say this to us as mental health professionals, judgment doesn't disappear. It's not replaced, but the metacognitive posture, that felt sense of I need to stay curious here, I don't know yet, can become

less available when the environment delivers something that feels complete, fluent, or coherent.

Next slide. Building on that, it's not just memory. There's really a second step the research stream points to. When information is readily available outside of the mind, people don't just retrieve it later. They tend to encode it less deeply in the first place because the brain is economizing. That's a well-established pattern in cognitive offloading work when the environment becomes a trusted storage layer, internal rehearsal, and that internal checking often decreases. We don't do it out of negligence, but it's because the task has been redistributed.

When we look at that in psychotherapy, if I know a coherent summary or clean note is waiting for me, my mind has less reason to do that slow, usually invisible work that turns an hour into the clinical understanding, the quiet replay of the emotional arc, the tracking of the subtle shifts in the narrative, and the building of that memory trace that holds not just the facts but the meaning behind them.

The shift is not simply in recall. It's actually recall equals reliance or leads to reliance. When reliance increases, something else becomes important, which is framing. What AI actually is giving you isn't a raw transcript. It's an interpretation. AI is designed to give you a structured narrative. Again, I'm going to keep using this word because I think it's important, but it gives a set of highlighted salient elements.

We know from decision support research that when a system provides a plausible frame to work off of, humans often anchor to it, especially under time pressure or cognitive load. Sometimes this reduces our own independent information seeking and verification. It's one of the mechanisms that is really described in automation bias research a lot. That's why the real, again, clinical question isn't, is the summary accurate? It's, how does the presence of a fluent, coherent frame shape our downstream judgment? What we look for in the next session, maybe? What hypotheses feel more available, and what uncertainty we tolerate?

This connects directly to metacognition. If the system's output is coherent and confident, it can feel truer, and that can really shift our confidence calibration if something feels truer. Work on explainability and metacognition highlights that how information is presented can change what people think they know, not just what they know.

Again, I don't want to treat this as fear-mongering. I think it's a predictive cognitive trade that we see. External access increases convenience, convenience increases reliance, and reliance increases the power of framing, sometimes quietly. The design or use implication here is practical. We want workflows using technology that preserve an internal first pass of meaning-making so the AI output becomes a second reader, not the first author of the clinical story. We'll talk about that in a little bit.

Next slide. Up to now, we've talked about memory and monitoring, but I think this is where it gets a little clinically intimate because in our work, psychotherapy isn't just about recall. It's about sense-making. Sense-making is a messy, iterative process of turning really complex, ambiguous human material into meaning formulation. Some of the things we've already talked about are hypotheses, the rupture, repair, and

relationships, risk signals, or like transference, counter-transference dynamics, all of that shifts. It's not just reduced cleanly to topics discussed sometimes in the work that we do.

The human AI sense-making research helps name what changes when AI enters this scene. When a system highlights themes or summarizes those moments or offers the interpretations, those outputs can often function as the interpretive anchors that we're using, which that starting frame, that subsequent thinking tends to organize around. Anchoring is really a normal cognitive phenomenon.

The first coherent interpretation carries disproportionate weight, especially when it's fluent and plausible, because AI is designed to do that. Decision support, research has long shown really that reliable systems can become a default frame for judgment even when users remain well-intentioned and thoughtful. I think that's important for us to know as therapists, because we are thoughtful and we are intentional about the work that we do.

What happens next is subtle. I do think that clinicians still engage in sense-making, and that's important, but the shape of what that sense-making looks like can shift. Instead of generating multiple competing hypotheses or exploring those tensions, our exploration can narrow, and that work becomes more confirmatory than generative unless we actively intervene.

I think this is where the slide title really matters. Meaning really does become co-constructed, but not in the traditional psychotherapy sense of clinician-client co-construction, but it becomes more of a distributed cognition sense. The system is now a participant in that interpretive environment. We have to be aware of that.

Von Zahn and colleagues work on explainable AI, and metacognition is really relevant here because it emphasizes that the way a system presents reasoning, its confidence, its clarity, or explanatory style, that can actually really influence a person's sense of what they know and how certain they feel. In clinical work, that matters, obviously, because confidence calibration is part of safety, knowing what we know and knowing what we don't yet know.

The point of this slide is not, again, AI is bad at meaning-making. It's AI changes the ecology of meaning-making. If the tool is supplying that early frame for us, we have to protect our capacity to stay spacious, hypothesis-rich, and really tolerate ambiguity. Otherwise, the range of interpretations we genuinely consider is shrinking without us really noticing.

Next slide. One of the promises of AI, and honestly, one of the reasons many clinicians are adopting it, is speed. Notes are drafted faster, patterns are surfaced sooner, the administrative loop closes more cleanly. I know we've circled this earlier, but it matters enough to name directly. Clinicians aren't practicing in spacious conditions anymore.

We actually are in a field where productivity, expectations, the documentation demands, the back-to-back schedules, they already squeeze reflection. Not because we aren't valuing reflection, but because the system doesn't reliably make room for it. AI isn't arriving in a neutral setting. It's entering in a workflow that already rewards

completion and fast. This is where automation bias and metacognition findings really matter.

Across studies that have been done on automated decision support, a consistent pattern shows up, that under time pressure and higher workload, people monitor automated outputs less closely and more likely to accept the suggestions with less independent verification. Again, not because we don't care, [laughs] but because monitoring takes cognitive capacity and the environment is pushing for speed.

The automation bias literature frames that as an adaptive reallocation of attention. When cognitive resources are strained, people rely more on the decision support to compensate. If it's reliable, it can be performance-enhancing. If it's not, or if it frames the task too narrowly, it can actually introduce new errors or blind spots. When we now layer in the concept of metacognition, the research summarized in the deck here points to another important distinction.

When guidance is readily available, confidence can rise without a parallel increase in insight. In other words, this feels handled. It's not the same thing, as this is deeply understood. In our work in psychotherapy, we know that depth often comes from slowing down, sitting with that uncertainty, noticing what doesn't quite fit. It's also resisting the urge to close the loop too quickly. Efficiency doesn't remove those capacities, but it can reshape when they show up and whether the workflow invites them.

With AI, it's not that AI makes clinicians really stop thinking. It's actually more humane, maybe more accurate to say that AI can become the default cognitive shortcut inside a system that's already starving clinicians of time to think. When that happens, reflection can drift toward the convergent, confirming what has already been surfaced rather than the divergent work psychotherapy often requires. Again, like we talked about, generating some alternatives, holding different hypotheses, staying open longer.

I think the takeaway of this slide is simple, but it's consequential. Faster isn't the same as deeper. If we want clinical judgment and really formulation to remain central to the work that we do, we can't rely just on clinician willpower alone. I think we do need workflows that intentionally protect that metacognitive checking, especially under the very conditions where speed is most tempting, and monitoring naturally declines.

Next slide. I want to pause here and just pull together again what the research across these areas is really pointing to, and maybe just to reiterate these concepts along the way. When you lay cognitive offloading and automation bias and metacognition side by side, we see a pattern emerge. Supportive systems really tend, again, to reduce the effort. They make tasks feel easier. They increase our confidence. That appeal isn't hypothetical.

When our capacity is thin, and let's face it, it's thin every day, and it gets even thinner, any relief can be protective rather than indulgent. The same body of research that we're talking about here also flags a parallel effect. When support is reliable, readily available, again, we've talked about some of these concepts as a repeat anchor, people tend to monitor their own thinking less closely. In other words,

that metacognitive oversight, the checking, the questioning, and the early interpretation, that declines because the environment has made it rational to do so or to conserve that effort.

The key nuance for us, this doesn't typically show up as an obvious mistake. It actually shows up as fluency. Decisions feel easier to us. The narrative feels coherent sooner. That coherent sooner is where I think curiosity is traded for that sense of closure. In the sense-making research, once there's a plausible frame on the table, exploration can narrow unless something actively reopens it.

With AI or technology, the frame is often delivered early and fluently. Everything that it gives you, the storyline, the structure, the suggested diagnosis, the default drift can be toward confirmation of those things rather than that continued discovery. Again, there's a clinical trade-off that I want to hold in this. AI does restore capacity. We know this. There's lots of research we've done internally at Upheal that shows this. It gives us time, attention, and emotional energy.

Unless we design and use it intentionally, the capacity can really be spent more on throughput rather than on more reflection. In psychotherapy, I think curiosity isn't a nice-to-have. It's actually a safety and effectiveness variable for us. It's actually what we do. It's what protects against all of the premature formulation that we're talking about. In oversimplifying things like case conceptualizations, and that slow drift is, I already know what this is. I don't want us to reject the capacity gained by using AI. I want us to protect the curiosity as clinical capacity by building in prompts and pauses that keep exploration alive, even when everything is easier and faster. Next slide, please. We've already been circling a few versions of this idea. Again, repeatable anchors, offloading closure, what feels complete, and the way that tools shake the attention. Here's what I want to add on this slide. A clearer description of how the shift actually shows up in real time and why it can be so easy to miss. The research points to something that clinicians often overlook precisely because it doesn't register as a mistake.

With that reliable automation, the most cognitive change is a quiet change in our monitoring. The system, again, starts to feel good enough, and the internal impulse to pause or widen our frame or do a second pass becomes less automatic. What makes this clinically relevant is that the signal itself usually isn't error, it's ease. Decisions, again, are going to feel smoother because conclusions are cleaner. The work is taking less effort, and ease can be persuasive. Our minds treat that fluency that AI gives us as a cue that we understand.

The risk here, obviously, is that ease can be mistaken for improved judgment when it's really just reduced friction in the thinking loop. Again, that doesn't mean the output is wrong. I want to be clear about that because AI is very sophisticated in that. It means that the workflow can change the timing of the reflection when the clinician's second look gets recruited. When I say that thinking quietly changes, I'm talking about that texture of clinical reasoning or the overlay shifts over time. There's fewer moments where we're naturally reopening the hour, fewer internal wait, like, "What am I missing here?"

Those interruptions, less spontaneous self-checking, especially when the output arrives within 90 seconds, [chuckles] coherent, all complete. That's really what I

mean by intuitive drift. It's a gradual change in how curiosity or questioning and that self-monitoring really show up when we have an external system carry more of the weight for us of that first pass. As you think about tools that you use, if it's AI or otherwise, pay attention to where ease really shows up. When something feels easy, ask yourself whether your monitoring is still online.

A good example that I thought about for myself outside of even just AI, when I was in grad school, I often found that the only way-- I told myself the only way that I could memorize the material for any exam was to write everything down, write until your hand cramped, repetition, repetition until it was memorized. I often, in that process, was not always examining whether or not I understood the material, just that I had it memorized. That's a very similar thing that we see here. I think that's a bridge to where we're going next is to understand that drift and how to design and practice against it, because we have to talk about metacognition, how we know what we know, and how we notice what we don't know.

Next slide. Metacognition, I've used that term a little bit, is how we monitor our own thinking in real time, how certain we feel about things, how much doubt we're actually willing to hold onto, whether we decide to pause, test things, or seek another mind, consultation, or supervision, another colleague, however that looks for you. It's really the internal process that tells you, "This feels clear, but I might be missing something." A key point in the research stream emphasizes that metacognitive monitoring isn't automatic. It does cost us effort, attention, and cognitive space. It's something we do actively, not just something that runs in the background.

I think that's exactly why reliable automation changes things for us. If you look at some of the studies, across the studies of humans interacting with automated systems, there's a consistent pattern that shows up. Again, when reliable guidance is available, people engage in less self-monitoring. They check their own reasoning less often. They rely more on the output because, again, the environment is trustworthy enough; it's less necessary, and something just feels handled for us. With respect to relevance in psychotherapy, I think clinical judgment really depends heavily on our ability to notice uncertainty in our thinking and to treat uncertainty as an instruction to stay curious rather than move forward.

If technology, if AI is drafting a narrative for us, the risk isn't really just accuracy. It's that the metacognitive prompt, that felt signal that says, slow down, interpret, interrogate, whatever you want to call it, gets quieter. It doesn't mean it goes away. It gets quieter. I think this sets up the next point for us. One of the clearest findings across the research is that confidence and understanding can decouple meaning. As I've talked about in previous slides, that workflow, really, because it feels clearer even when insight hasn't deepened. That's where we're going to go next.

Next slide. Again, like I said, one of the clearest findings across the automation and metacognitive research is confidence and understanding don't always move together. When the systems are providing the guidance, people often feel more confident, even when their actual understanding really hasn't deepened. The literature describes that as a decoupling between confidence and insight. What I want to be careful about in how we frame this, especially with us as mental health professionals, it's not about laziness, carelessness, or outsourcing responsibility. The

research explicitly frames it as predictable response to external support that appears reliable and reasonable.

The mechanism, really in plain terms, is, again, when guidance is there, when it's available, it's almost like you have a clinical supervisor there all the time saying yes or no. The internal signal, though, to slow down and double-check in ourselves or trust ourselves more, weakens. Monitoring declines, but certainty still rises. Again, if we look at that in psychotherapy, where the goal isn't just correct content, it's accurate sensemaking, that decoupling can show up clinically as a quicker formulation that feels obvious sooner than it should. It feels like earlier closure around what the case is.

It's less tolerance for ambiguity. It's fewer moments of counter formulation, that deliberate practice of asking what else could this be, that differential piece of what we do. Again, the point of this slide is not that AI will make us overconfident. It's actually that AI can make the work feel clearer sooner. When we're under real clinical constraints, that felt clarity that we're provided from that becomes that persuasive. The workflow rewards the forward motion, and the mind naturally follows. This is why the distinction at the bottom of the slide matters. I think the research consistently frames this as a monitoring issue, not a competence issue.

If it is a monitoring issue, it's not solved by trying harder. It's solved by creating deliberate moments where clinicians are invited by design or practice to re-engage in that metacognitive process before the narrative hardens for us. The hardest part is that these shifts don't often feel like risk. They feel like clarity. Again, clarity is persuasive. That's why automation bias can be invisible somewhat from the inside. Next slide. One of the most challenging aspects of what we've been calling intuitive drift is that it's hard to detect from the inside. Automation bias research points out that reliance on an automated system often doesn't feel like reliance. It feels like competence.

The system works. The output of the system makes sense. The workflow feels efficient. Nothing in your experience tells you, "Be careful, your monitoring might have changed." The sense-making findings add the next layer. When a system, again, provides those early frames, those outputs become the anchors. You still think, you still reflect, but again, like I said earlier, your reflection starts from within the frame that's been provided to you. That reliably narrows exploration unless something actively reopens it for you. Part of the work that we do really does depend on resisting that early clarity and letting the meaning unfold across more sessions.

Sometimes, again, the pressures of the environment don't allow for that. Sometimes we're like, "Hey, you need to give us a diagnosis immediately, and you need to be billing, you need to do this." Some of that isn't really allowable on environments, but the true work of psychotherapy does allow for meaning to unfold across time. The risk isn't an error event. It's that the interpretive range quietly contracts without triggering any internal alarm because our experience objectively is improved by fluency and not confusion. This slide is really a warning about what we won't notice unless we're intentional.

We won't notice monitoring declining because things are still working. We won't notice the anchoring that's taking place because the first frame will feel reasonable to

us. We won't notice fewer alternative hypotheses surfacing because nothing obviously breaks. Where I want to be clear, these shifts, again, are not about clinician morality or motivation. They're about predictable cognitive adaptation to the environment, which is exactly what the next slide names. Next slide. This slide comes out of a sense of some of the things I'm seeing in a lot of the just conversations between therapists. I want to be clear here, because this is where conversations can go off and you go off track, clinician to clinician.

When we talk about risks like automation, bias, anchoring, or intuitive drift, it's really tempting to frame this as a character problem. We often hear things like, clinicians should just be more careful. We should try harder to think critically. Again, I actually see this all the time on clinician forums, and the framing of statements like, "Why use AI at all? We've done just fine for decades, doing notes the old-fashioned way." There are a lot of characterizations that choosing to use AI or technology somehow an ethical and a moral failure on the part of clinicians, but again, the research doesn't support that framing. What it supports is something more uncomfortable, but maybe more compassionate.

Human cognition adapts to the environments it's working inside, and it's not as a moral choice, but as a functional response. When the system reliably reduces effort, the brain treats that as a cue that this is safe enough. I don't need to spend as much time monitoring energy here. The impact of moral injury and working inside a broken system tends to move clinicians more towards choices, to reduce the cognitive load, to stay in the field, and to try to keep doing the work that they love. The part that I want to juxtapose with the reality of clinical work, again, when we're introducing a tool that increases speed and reduces friction, it's not surprising that monitoring gets spent elsewhere.

That's not clinician disengagement. I think it's resource allocation amidst the environment. The risk emerges, really, from context, not the intent. Not because, again, we don't care as clinicians, or because we're less ethical, but because, again, the environment and the tool chain make fast and finished feel like the correct defaults, and a lot of the systems that we work in. I think this leads to a practical implication. If monitoring declines, unless something actively prompts reflection, then the question isn't, how do we get clinicians to try harder? We really need to think about do our tools and workflows that we're engaging in actively support metacognitive engagement, especially when we have real-world constraints.

Again, consider that when you're looking at different AI tools to support you, because if they don't, I think that we continue to misdiagnose a structural cognitive phenomenon as an individual moral one. If you've ever felt your thinking get tighter at the end of a long clinical day, or less exploratory or more good enough, that's not a values problem. That's cognition under constraint.

AI can either protect reflective depth or compress it even further; the choice is really yours. Once we stop moralizing this, we can finally do something useful with it, because again, what we're seeing these effects are predictable, which means we can move from predictable risk to intentional practice with the use of it. Up to this point, we've been really deliberate about naming the risk, and I want to be equally clear about what the literature is not saying. The research does not suggest that

automation inevitably undermines judgment, or that clinicians become dependent or less ethical simply by using it. Again, I think we framed that.

What it does show across cognitive offloading, automation, bias, and metacognition is something more practical. The cognitive effects of automation are predictable, which means they're not just risks to warn about. They're actually patterns we can design around, and something we actively do in this field, if we're in positions where we have a design approach to AI. For example, Gilbert's work on offloading really helps explain why reliable external supports reduce the need for internal rehearsal and monitoring. That's not pathology. It's the adaptation we keep talking about, but it tells us exactly where to intervene.

If monitoring fades by default, then you have to build monitoring back in on purpose through workflow prompts and norms in your system design. von Zahn's work on explainability and metacognition helps name that another layer, the way that information is presented changes confidence calibration. If a system's output is fluent and complete, it can conflate that certainty we've talked about. The fix isn't just be careful. That's something that initial conversations around AI we would talk about, but the fix is really metacognitive scaffolding using structures that keep the clinician actively testing, not just accepting.

What does intentional practice look like for us using tools with psychotherapy? It's not really complicated, but it is a deliberate choice. What it looks like is a first pass before second pass. What you want to do is you want to capture your impression of the session, your felt sense of what occurred, your provisional formulation, whatever risk you read before you even look at the AI summary, and then you want to compare it. That preserves your clinician-generated salience before you look at an AI-generated output. You want to engage in counter formulation on purpose, and one forced question might look like, "What else could this be? What could, what would I think if I didn't have the summary in front of me?"

That keeps hypothesis generation alive. Engage in a confidence check. When something feels very clear, maybe add a pause. What makes me feel confident in this? What would lower my confidence? This really represents that metacognition in action, that calibration, not paranoia. Again, we want to use AI as a second reader, not the author. What I would actually say is, treat the output like a colleague's draft, useful, but not definitive. Again, the big idea is simple. If automation predictively nudges us towards ease and closure, then we respond by designing for reflection and oversight, especially under those real-world constraints we're talking about.

Once you start thinking this way, the next question becomes, again, what specific practices and design choices keep the metacognition alive without adding a ton of burden back. Next slide. I want to name a false choice that shows up a lot in clinical environments, that you can be compliant or you can be curious, but you can't be both. Compliance is real. It's how systems manage accountability that shows up as medical necessity, risk, continuity, all leading to reimbursement. Curiosity is also real. It's how psychotherapy stays alive and how we track our ambivalence, notice what doesn't fit. That toleration of uncertainty we've been talking about.

Again, the problem is that documentation systems don't just record our thinking. They start to train us how to think. The dominant job becomes, produce a defensible

narrative. When that becomes that way, the clinician's attention gets shaped towards what the template can hold, the symptoms, the categories, the interventions, the progress markers. That structure is really helpful until it becomes the primary lens through which we're listening because, really, the core clinical stance we need in the room is often the opposite of what the record rewards.

In the room, a mental health professional might be tracking, what's the function of this silence? Why is this affect shift happening right now? What's the relational bit happening under this? The compliance-oriented mind is tracking: What do I need to show? What looks coherent? What will withstand scrutiny? What will prevent clawbacks? That's the tension I've often found myself in my own private practice, and I think one that often leads to paralysis in my documentation or many others. The key point, that shift isn't about clinicians becoming less skilled; it's about the finite cognitive, emotional resources that are redirected under this pressure.

The empathy literature is pretty blunt about this. Empathic engagement isn't free. It does require emotional and motivational investment, something a human clinician is actually spending in the relationship. That investment is part of what makes care feel real and specific, and it's part of what supports outcomes, but it also means empathy is effortful and it can be depleted or narrowed when the systems pull our attention elsewhere. When the environment keeps asking for standardized compliance-ready outputs, it can unintentionally crowd out the conditions that protect that empathy or that exploratory stance.

Now, when we layer in AI, AI absolutely can help here. If it's used to offload the clerical weight, to protect that therapist bandwidth for presence, the formulation, and the relational work that we do. I've often talked about this in a lot of my other talks, but there's a very real risk. If AI gets deployed primarily to accelerate documentation throughput, it can actually reinforce the idea that the successful clinician is the one who is efficient. Again, research cautions us with this. Over-reliance can make care feel more mechanistic, and it can also create that automation bias that we're talking about, where human judgment gets quietly overshadowed by the output.

The clinical question I want to leave hanging on the slide is not again, are we compliant? We have to be. It's are the tools that we're using helping us stay compliant without training our minds away from curiosity? The moment compliance becomes the main cognitive task, curiosity becomes the first casualty, if we're honest. Psychotherapy really becomes less psychologically rich, even if the note looks perfect. Next slide. Again, when AI is introduced into documentation-heavy environments, it isn't arriving as a neutral add-on. It usually is getting deployed to solve the very real, specific organizational problem of make documentation faster, cleaner, more consistent.

In my work, this is what enterprise customers might be seeking. They often do so in a way to also help the clinicians, but it is also to manage the risk from compliance. To be fair, again, these systems often succeed at that. They can accelerate the workflows. They can reduce the administrative strain, which is a real need in psychotherapy services operating under load. The tension that the research points out is that, again, speed and consistently unintentionally reward faster closure. When the workflow is optimized around completion, users of the product naturally become more likely to accept the outputs that align with the institutional priorities.

Again, not because they're disengaged, but because that's the system that's designed to move things forward efficiently and get them their money. The risk isn't automation replaces clinical thinking. It's actually automation changes when reflection happens. If reflection isn't explicitly supported, it becomes easier for curiosity to become compressed even further. That's why the slide uses the phrase reinforced attention. Without intentional design, automation can unintentionally strengthen the compliance side of the equation at the expense of the reflective clinical thinking, not by taking the judgment away, but by quietly reorganizing the workflow around that speed.

Next slide. I think this is the point where research becomes pretty actionable for us. Everything that we've discussed, offloading, automation, bias, the metacognitive drift, it converges on one practical conclusion. Reflection does not reliably survive efficiency environments by default. When systems reduce effort and increase the speed, monitoring doesn't just stay on; it tends to fade unless something actively supports it. That's why scaffolding matters, not as a moral correction that clinicians should try harder, but as a cognitive support that matches how human judgment actually functions under load.

I want to point out something here. Again, the research isn't suggesting that we need to choose between compliance and curiosity, like I said before. It actually suggests if we want both, reflection has to be designed and practiced intentionally. Again, scaffolding in this context means creating small but meaningful structures that prompt comparison before commitment, keep formulation, maybe that provisional beat a little longer, and invite the self-questioning at the moment it matters before that narrative hardens into the story of the session. In other words, scaffolding doesn't replace thinking; it preserves it.

If you take one thing from this slide, it's this: that compliance and curiosity, again, don't have to be opposites, but they won't coexist by accident, especially in systems optimized for the completion. Now that we've named scaffolding as the lever, the next question is, what does this actually look like in practice, and how do we design for that ethical fidelity and that clinical growth without the burden being added back in? Next slide. Sorry, I'm trying to go a little bit quicker here because I know I've got a lot more. If we take this research seriously, the ethical use of AI isn't only about consent, privacy, or compliance. Those are very important.

It's also about whether we're preserving the clinician's capacity to think or reflect and exercise judgment over time. Again, workflow design, not neutral. Systems aren't just storing this information; they shape what we attend to and what feels important, what feels complete. I think that that feels worth revisiting. In other words, workflows are quietly teaching clinicians what matters through what they surface, what they hide, and what they reward. This is where the earlier threads connect.

Again, when systems are optimizing primarily, when I say systems, the conditions we're working in, the environments, when they're optimizing primarily for what's billable, defensible, or standardized, they unintentionally reinforce the narrow cognitive loop of closure, document complete, good enough, close the note within 24 hours. When systems also scaffold reflection, prompting comparison, uncertainty, and engaging in alternative hypotheses, they make space for clinical growth, not just compliance.

When we're thinking about design, the very practical nature is when we want to build workflows that help clinicians stay metacognitively engaged, that question becomes, are they? Are we doing that, or does it quietly bypass that engagement by making it just feel finished? Again, ethical fidelity isn't about only the data we use or how it's protected. It also includes protecting that process that really helps define our clinical expertise. All of the things we've talked about that leads to reflective judgment, especially under the real clinician constraints that we're currently surviving in.

Next slide. Again, making this concrete because scaffolding can sound abstract until you really can picture it in a real clinical day. The simplest way to look at that is AI outputs should be inputs, not conclusions. AI can be useful for organizing content, drafting the structure, and surfacing the themes, but the moment we treat the output as the answer, we risk letting it become the default frame for how we really understood the session in the first place. When that happens, our reflection again narrows. We end up refining what's been presented instead of generating and testing against our own hypothesis here.

What it looks like to keep that clinical judgment central is, again, to do that first clinician pass before you look at the AI. The 30 seconds, ask yourself, what were the one to two moments that felt really clinically hot or significant to me? What felt unresolved in the session? What is my working hypothesis? What is the alternative explanation I need to keep on the table? Then you look at the AI output and use it in the way that, again, like we talked about earlier, that it's authoritative. Put the reflection step back before the note becomes final. Then keep judgment clinician-led. Those are some things I said earlier, but I think it's important to make sure that we put that at the forefront of our minds.

Next slide. I want to close by coming back to what sits underneath all of this, which is cognitive responsibility. AI isn't changing necessarily what clinicians do. It doesn't just do that. It does change how we attend, remember, reason, or make the meaning out of the work we do. Again, the research we've been reviewing is consistent on a core point. When automation becomes part of our decision-making, our brain will adapt. Our cognition, our thinking, will adapt. Our monitoring will shift. Our confidence will rise. Sensemaking reorganizes itself around outputs. None of that, again, is inherently negative, but none of it is neutral either.

Again, here's the piece I want to name for us. This is happening inside environments that already strain attention. When we have high workloads, tasks are complex, time is tight, the reliance on the supports around us, they are an adaptive strategy. Our attention is getting reallocated. This is all true. It is attention. It's how attention does work under pressure. In many clinical settings, pressure is the default condition for a lot of us. It's a default condition, not an exception. When we talk about responsibility, I'm not just talking about whether we choose a compliant tool.

I'm talking about whether we remain aware of that the predictable cognitive shift that we know that these tools create, and whether we're engaging with tools that have workflows that make room for that reflective engagement. The risk that we're talking about is not about the clinician stopping thinking. It's actually understanding where the starting point of thinking shifts. The automation literature, again, is not just describing risk factors like workload and time pressure. It's also pointing to mitigators.

That's the stuff that we're designing around, which is making the reasoning more visible, the strengthening accountability, and then designing outputs in ways that support appropriate reliance instead of default acceptance. The cognitive responsibility includes how we use tools, but also includes how we design the workflow around them. Especially in crunch conditions where monitoring doesn't persist on its own. From this perspective, I think ethical responsibility isn't just about choosing the right tools. It's really about taking responsibility for how those tools shape our thinking over time.

Next slide. One of the things that I really hope is clear by now, again, adaptation isn't a problem. It's a feature of our cognition. We are adapting to tools, environments, and systems all the time. That's part of what makes expertise really possible in the first place. What's truly different with AI is that the adaptation can happen quickly. It often is below our conscious awareness. That's one of the things I've really tried to bring up to the surface today. That kind of awareness really that we need to be thinking about is that metacognitive awareness. It's noticing what am I accepting too quickly? What am I no longer scanning for?

What did I use to generate internally that I'm now receiving externally? Those sort of things. The point I want to land here, again, is that research suggests this is not an optional add-on to clinical skill. It is becoming a core competency. I think one that becomes much more important, not less, as AI becomes more integrated into the care that we do. Again, the goal isn't to resist change. It's to stay oriented to what's happening inside our own cognition as the workflow changes around us and to build small, repeatable structures that protect that reflective depth, even when everything else has us cruising toward autopilot.

We can probably skip the next couple slides to take questions at this point, because I realize that they're probably more just key takeaways we've talked about. I really appreciate everybody's hanging on, because I know we're nearing the end of our time.

Jody: Thank you, Ted. We are at the top of the hour, and I don't know that we will have time to take questions today, but we are going to put your email in the chat so that people could reach out to you directly. I want to thank you for your participation. Again, to everyone who submitted their questions in the Q&A, we're going to put Ted's email in the chat. Ted, I want to thank you for joining us, and thank all of our listeners for joining us today, too. We'd also like to thank Upheal for making this webinar possible. As mentioned at the beginning, a recording will go out in about two weeks. That will also include today's presentation slides.

For those that joined us today, if you could take the survey immediately following the webinar, that would be fantastic. Thank you again. Thank you, Upheal. Thank you, Ted. I hope everyone has a great afternoon. Take care.

Ted: Thank you.

[01:00:28] [END OF AUDIO]