Why Most AI Chatbots Fail to Convince Us (And How to Fix Them)

By Dr. James L. Norrie, DPM, LL.M | September 17, 2025

Have you ever asked a chatbot for advice, only to feel as if it were parroting back a script? You're not alone. Our research shows that most AI chatbots come across as flat, impersonal, and honestly, quite forgettable. They mimic conversation but rarely persuade us to act. The missing ingredient isn't simply better facts, it's personalization.

And by that I don't mean the superficial tweaks of giving a chatbot a male or female voice, or swapping in a British accent for an American one. Too easy. Those cosmetic flourishes are little more than window dressing. The real challenge is deeper: how to move beyond rote, generic responses to create human trust in chatbots that persuade? And make sure they are used in ethical and responsible ways that make lives better?

The Human Problem with Machines That "Talk"

Think about the last time you brushed off an AI assistant's suggestion. Maybe it nudged you to update your password or warned you to double-check a suspicious link. Chances are you ignored it—not because the advice was wrong, but because it was delivered in the same bland, one-size-fits-all tone. Humans don't just want information; we want connection. And connection requires trust, a sense that the person—or machine—on the other end actually understands us. It's simple but profound: we need to feel heard before we are willing to listen. Conversely, those with various mental health conditions can be unduly influenced by online sources, and we need extra caution in that case.

Adding Genuine Personality to AI

That's where psychology comes in. Decades of research confirm what most of us already know intuitively: personality shapes how we respond to rules, risks, and persuasion. In our work, we explored this through the lens of cybersecurity, studying how people make choices in online scenarios where danger often hides in plain sight. Some individuals are natural risk-takers, others crave structure and predictability. Some defer readily to authority, while others instinctively resist it. These differences are not superficial preferences; they are rooted in enduring personality traits that shape perception, appraisal, and action. Similarly, mental health conditions can create undue vulnerability to online influence that needs to also be embedded into AI chatbots to avoid harm.

When a chatbot speaks to every user in the same generic way, it fails to establish a psychological precondition that matters most in human interactions: trust. Without trust, even accurate information struggles to gain traction with most of us. The result is not simply reduced engagement, but an absence of influence at precisely the moment

when behavioral guidance is most needed. In psychological terms, the message lacks contextual congruence with the user's disposition. So it doesn't resonate deeply enough to alter decision-making in real time. In other words, without true personalization, persuasion falters exactly at the critical inflection point when decisions are made.

Our expert team set out to tackle this challenge head-on, and the result was a system that embeds genuine personality into AI interactions. Building on our patented **myQ™ instrument**—which maps an individual's risk tolerance and rule-following tendencies—giving previously generic AI the improved ability to adjust its tone, style, and phrasing to align with each user's specific and known personality traits and cognitive biases.

The difference is profound, and with many potential applications where AI chatbots or agents are being introduced. Instead of forcing users to select from a few generic "types" or voice options, the system adapts automatically in real-time, tailoring every AI response in subtle but powerful ways that increases the likelihood you will not only hear the advice, but act on it in the moment. We call this fusion of psychology and artificial intelligence **AIQ™**.

Getting the Facts Right, Too

Of course, personality alone isn't enough. An AI can be warm and funny, but if it gives you bad information, you'll stop trusting it. And trust is the very essence of influence, a necessary pre-condition to achieving influence with humans. That's why our system also uses an existing technique called **RAG**—think of it as fact-checking on demand. It pulls from curated, reliable sources before answering, reducing the "hallucinations" that plague many chatbots. The result is advice that's not just engaging, but also accurate.

Why This Matters

In many situations, people hesitate to ask questions when they don't know something. That hesitation stems from different sources—fear of embarrassment, lack of confidence, or concern about being judged—but it often stifles curiosity. Think of the average employee: how likely are they to call the security team or IT helpdesk to ask about something that might be dismissed as "basic" or routine? In that context, the fear of judgment becomes a powerful inhibitor.

So we tested a different approach. Instead of relying on employees to seek out human experts, we introduced a company AI chatbot to guide them through cybersecurity risks—helping them spot phishing scams and other online traps in real time. The results were striking. Within months, failure rates dropped by more than half, as employees began turning to the chatbot as their first line of defense at the very moments when risk was highest. Psychologists would call this a **pattern interrupt**: an intervention that disrupts automatic, often unsafe behaviors and replaces them with more deliberate choices.

And it worked. People didn't just consult the AI; they trusted it enough to act on its advice.

The implications go far beyond cybersecurity. Imagine a healthcare chatbot that reminds patients to take their medication in a way that feels natural to them. Or a financial advisor bot that nudges someone away from a risky investment using language that resonates with their decision-making style. In each case, the principle is the same: AI trust built through personalization turns information into better human choices.

The Bigger Picture

Presently, we're only scratching the surface of what hyper-personalized AI actually is and what it can do. But like AI itself, we are learning quickly. Done right, our method to style-align AI chatbots can become a powerful tool for good: improving safety, health, and financial well-being by speaking to us in ways that feel more personal and trustworthy. In fact, our pilot research showed remarkable improvements in orders of magnitude that were initially

But there's a flip side too. Any technology that can influence behavior this effectively must be used ethically. Without clear guardrails, the same tools could be used to manipulate, deceive, or exploit.

The lesson is clear: if AI is going to play a bigger role in our lives, it must learn not just to *inform* us, but to *connect* with us. Machines that talk like us will only truly matter when they also understand us. And we are rapidly learning how to do that, and how to embed guardrails to protect those among us who are most vulnerable. Contact us if you would like more information or assistance in personalizing your own AI implementations.

Author Bio:
Dr. James L. Norrie is a professor of Law and Cybersecurity, and Founding Dean of the Graham School of Business, at York College of Pennsylvania (http://www.ycp.edu). He is the author of *Beyond the Code: AI's Promise, Peril and Possibility for Humanity* (Kendall, Hunt 2025). Learn more about our free community of interest in ethical AI at: www.techellect.com or visit www.cyberconIQ.com to learn more about AI tools to keep your employees safer online. To purchase his book, click on the QR code, or visit: https://he.kendallhunt.com/product/beyond-code-ais-promise-peril-and-possibility-humanity