




© Irissca-AdobeStock

LIGHTS - Human/Well being

The psychology of fragile choices: consumer vulnerability in the age of generative AI

ESCP Impact Paper No.2025-27-EN

Isabella MAGGIONI
ESCP Business School

[The psychology of fragile choices: consumer vulnerability in the age of generative AI

Isabella Maggioni*

ESCP Business School

Abstract

As generative AI becomes increasingly embedded in everyday consumer experiences, its influence extends beyond utility into psychologically sensitive terrain. This paper explores how AI systems engage with consumers during moments of emotional or cognitive fragility, transforming from neutral tools into perceived companions that shape decision-making, identity, and behaviour. Drawing on recent behavioural and marketing research, we highlight the paradoxical role of AI as both comforter and amplifier of vulnerability, and introduce the HEART framework: a model that outlines five dimensions, Heuristic Dependence, Emotional State, Asymmetric Knowledge, Reliance Tendency, and Temporality of Vulnerability. The framework offers a new lens through which businesses can anticipate emotional risk, embed ethical safeguards, and design AI systems that protect user agency in moments of diminished resilience.

Keywords: consumer vulnerability, AI ethics, psychological safety, human-AI interactions

*Associate Professor of Marketing, ESCP Business School

ESCP Impact Papers are in draft form. This paper is circulated for the purposes of comment and discussion only. Hence, it does not preclude simultaneous or subsequent publication elsewhere. ESCP Impact Papers are not refereed. The form and content of papers are the responsibility of individual authors. ESCP Business School does not bear any responsibility for the views expressed in the articles. Copyright for the paper is held by the individual authors.

The psychology of fragile choices: consumer vulnerability in the age of generative AI

Introduction

The rise of generative AI has introduced powerful tools that are now deeply integrated in everyday consumer journeys, from health inquiries and financial decisions to emotional support. AI is no longer a back-end tool, it is a front-facing participant in emotionally charged decisions. Whether consumers are consulting AI for financial advice, mental health support, or product guidance, the interaction often occurs during moments of diminished cognitive or emotional capacity. Recent research shows that vulnerable users engage disproportionately with AI systems for guidance, comfort, or decision-making support (De Freitas *et al.*, 2023; Hermann *et al.*, 2024).

Marketing literature has increasingly addressed the concept of vulnerable consumers, defined as those whose personal, cognitive, or situational characteristics limit their ability to make fully autonomous and informed choices (Baker *et al.*, 2005; Hill, 2001). Such vulnerability is not determined by age or ability, but often emerges from context and emotional state, making AI-based touchpoints especially sensitive environments. Moreover, as Yap *et al.* (2021) suggest, technology can paradoxically serve both as a buffer and amplifier of vulnerability, especially during crises.

Historically, businesses have treated vulnerability as a fixed characteristic of certain demographics. However, vulnerability can arise episodically, affecting any consumer in response to personal, social, or economic stressors (Hermann *et al.*, 2024). For firms deploying AI, this reframing means that ethical responsibility cannot be confined to edge cases. Indeed, every AI touchpoint has the potential to engage a vulnerable individual. This interaction carries important implications for risk management, service design, and leadership strategy, all of which must adapt accordingly.

AI as a psychological mirror

Generative AI systems are increasingly acting as “psychological mirrors,” offering not only quick answers, but also emotional reassurance and behavioural nudges. While this responsiveness can help streamline decision-making, it introduces complex ethical concerns.

Empirical evidence shows that users often anthropomorphize AI, perceiving it as emotionally attuned or even wise, particularly during moments of vulnerability (De Freitas *et al.*, 2023). In such states, users tend to apply less critical scrutiny; they become more prone to overreliance, especially when AI mimics empathy or adopts a relational tone. What begins as a tool quickly becomes a cognitive shortcut, bypassing reflection and subtly influencing behaviour.

Without ethical guardrails, firms risk exploiting users rather than empowering them. Poorly timed nudges, emotionally manipulative outputs, or unfiltered advice can have real-world consequences, particularly in sensitive sectors like mental health, finance, and healthcare. In mental health, for instance, users experiencing loneliness or anxiety have reported turning to AI chatbots for emotional support, sometimes several times a day, and interpreting the responses as genuinely empathic (De Freitas *et al.*, 2024). While this might offer temporary comfort, it can also mask underlying needs or delay professional help. In

finance, similar risks arise when predictive personalization targets vulnerable users with high-risk financial products, nudging them toward quick credit, speculative investments, or services they may not fully understand. In healthcare, emotionally charged queries such as “Do I have cancer?” or “Should I ignore these symptoms?” can yield simplified or overly confident answers, lacking nuance or contextual safety. These outcomes are not merely design flaws, but they reveal a deeper structural problem: the disconnection between personalization and psychological safety.

Crucially, consumers are not just interacting with AI, they are indeed forming relationships with it. In emotionally depleted states, AI is experienced as calm, wise, and nonjudgmental. This perception reduces the emotional friction of decision-making, but also discourages critical evaluation. Behavioural science has long shown that stress, anxiety, and fatigue reduce cognitive capacity and increase reliance on heuristics (Backer *et al.*, 2005; Hill & Sharma, 2020). Generative AI, by its very nature, aligns with these mental shortcuts: it offers immediacy, certainty, and the comforting illusion of empathy. This creates a paradox: the more emotionally fragile users are, the more they turn to AI, and the more likely they are to be influenced by it (Hermann *et al.*, 2024).

This raises pressing questions about psychological dependency from AI. When a consumer repeatedly seeks AI affirmation during fragile moments, what happens to their sense of agency? Over time, reliance may shift from functional assistance to emotional validation delegating self-worth to algorithmic systems. As AI becomes embedded into daily life, firms must confront the challenge: how can they deliver value while avoiding emotional overreach and safeguarding user wellbeing?

The HEART framework: diagnosing and designing for consumer vulnerability in AI interactions

To help companies navigate this evolving terrain, we propose the HEART framework, a model that captures five dimensions through which consumer vulnerability emerges in AI-mediated contexts. Rather than treating vulnerability as a static trait limited to certain populations, the HEART framework recognizes it as a situational and behavioural phenomenon, shaped by emotional states, cognitive tendencies, and the very design of AI systems. Each of the five dimensions, Heuristic Dependence, Emotional State, Asymmetric Knowledge, Reliance Tendency, and Temporality of Vulnerability, offers a lens through which firms can better understand risk and reimagine their approach to AI design, communication, and governance. The HEART framework is a preliminary conceptual model that stems from ongoing qualitative research, including more than 30 in-depth interviews with consumers and analysis of online discussions in online forums and social media. These sources reveal consistent psychological patterns in how individuals turn to AI systems during moments of emotional or cognitive fragility.

1. **Heuristic Dependence:** In emotionally charged or cognitively demanding situations, users often fall back on heuristics to simplify complex decisions. Generative AI systems are built to deliver fast, fluent, and confident responses, which align perfectly with this cognitive tendency. During moments of stress or uncertainty, consumers are especially likely to accept these outputs at face value, bypassing critical thinking in favour of clarity and immediacy. When AI interfaces present suggestions with a confident tone and minimal friction, they effectively position themselves as low-effort decision aids. For companies, this means that even small design choices, like default recommendations, phrasing, or visual hierarchy can disproportionately shape user

outcomes. Designing for reflection rather than speed becomes essential in contexts where decisions carry emotional weight.

2. **Emotional State:** Feelings of anxiety, loneliness, grief, or insecurity not only increase the need for reassurance but also lower cognitive resistance to persuasive cues. In these moments, AI systems that simulate warmth, empathy, or attentiveness can feel especially comforting, sometimes more so than human interactions. However, this emotional resonance also carries risk. Consumers may interpret these affective responses as signs of care or understanding, even though they are generated by pattern recognition rather than true empathy. In vulnerable states, this misperception makes users more susceptible to suggestion, particularly if the AI affirms their concerns or offers seemingly tailored advice. For companies, acknowledging the emotional dynamics of interaction means moving beyond engagement metrics to consider the psychological impact of tone, timing, and framing. A chatbot that feels “kind” in everyday use may exert outsized influence during a crisis.
3. **Asymmetric Knowledge:** AI systems operate with an increasing degree of opacity. Most users have little understanding of how large language models generate responses, what data they use, or what optimization goals they serve. This asymmetry creates a fundamental power imbalance. When users lack understanding of how AI systems work, they tend to defer to them by default and mistake seamless and confident personalization for authority, rather than evaluating alternatives critically. The widening gap between system intelligence and user comprehension means that companies have a responsibility to embed transparency into interaction design. Not through dense disclosures, but through interface cues, feedback loops, and opportunities for users to question, understand, or opt out. Failing to do so not only erodes user agency but can also lead to legal, reputational, and ethical backlash.
4. **Reliance Tendency:** As AI tools become more accessible and intuitive, users may begin to turn to them habitually for emotional reassurance or decision validation. Over time, this pattern can shift from convenience to dependency. When a system is always available, nonjudgmental, and seemingly wise, it could become a proxy for reflection, feedback, or even self-worth. This phenomenon is especially pronounced in domains such as mental health or wellness, where users might consult AI tools repeatedly for affirmation or guidance. For organizations, this presents a delicate challenge. While consistency and availability are key drivers of user satisfaction, there is a tipping point where repeated engagement signals not utility, but emotional reliance. Monitoring usage patterns, offering human fallback options, or setting soft boundaries on certain forms of engagement can help avoid scenarios where users substitute AI for professional, social, or relational support.
5. **Temporality of Vulnerability:** Vulnerability fluctuates across time, context, and life events. A user who is highly competent in one moment may become emotionally fragile in another. These episodic shifts often go undetected by AI systems, which are not currently designed to sense or respond to changes in emotional or psychological state. This invisibility poses significant design and governance challenges. If a system treats all users as equally rational, equally informed, and equally resilient, it will inevitably fail those experiencing moments of distress. Recognizing that vulnerability is situational and time-bound calls for systems that are context-aware, adaptable, and designed with escalation pathways in mind. It also means that leadership teams need to rethink risk: not as a rare edge-case event, but as a structural feature of everyday engagement.

Conclusion

As generative AI becomes a constant presence in consumers' lives, its influence now shapes decisions, emotions, and perceptions of trust, often at moments of heightened vulnerability. This shift calls for a deeper sense of responsibility across the AI lifecycle, from design to leadership. The HEART framework offers a new lens for this challenge, urging companies to recognize users not as predictable decision-makers, but as emotionally complex individuals whose needs and agency fluctuate over time.

Key interventions for practitioners include auditing AI touchpoints for heuristic bias, adapting content to users' emotional states, embedding transparency features to reduce knowledge asymmetries, monitoring patterns of overreliance, and enabling context-sensitive responses to fluctuating vulnerability. Even subtle design changes, such as interface cues, soft boundaries, or escalation paths, can meaningfully enhance psychological safety. Beyond design, HEART can inform broader governance practices. Companies may integrate it into internal audits, ethical review protocols, or design team training. It can also support external-facing policies like AI disclosures or consent design. In highly sensitive sectors such as healthcare or finance, HEART can serve as a checklist to align AI systems with consumer protection and emotional wellbeing.

Incorporating HEART into AI development means asking difficult but necessary questions at every stage: How might this system interact with someone in distress? What shortcuts does it enable, and when are they risky? Are we designing for efficiency at the cost of autonomy? Are we nurturing trust or simply minimizing resistance? Ultimately, HEART invites a shift in mindset: from building systems that work perfectly under ideal conditions, to building systems that protect users when conditions are anything but ideal.

As AI continues to mirror, mediate, and even shape human experience, the question is no longer whether vulnerability will surface, but whether organizations will be ready to meet it with care.

References

- Baker S. M., Gentry, J. W., & Rittenburg, T. L. (2005). Building understanding of the domain of consumer vulnerability. *Journal of Macromarketing*, 25(2), 128–139.
- De Freitas, J., Uğuralp, A. K., Oğuz-Uğuralp, Z., & Puntoni, S. (2023). Chatbots and mental health: Insights into the safety of generative AI. *Journal of Consumer Psychology*, 34(4), 481–491.
- De Freitas, J., Oğuz-Uğuralp, Z., Uğuralp, A. K., & Puntoni, S. (2024). Why Most Resist AI Companions. *Harvard Business School Working Paper*.
- Hermann, E., Yalcin Williams, G., & Puntoni, S. (2024). Deploying Artificial Intelligence in Services to AID Vulnerable Consumers. *Journal of the Academy of Marketing Science*, 52, 1431–1451.
- Hill, R. P., & Sharma, E. (2020). Consumer vulnerability. *Journal of Consumer Psychology*, 30(3), 551–570.
- Yalcin, G., & Puntoni, S. (2023). How AI Affects Our Sense of Self - and Why It Matters for Business. *Harvard Business Review*, 101(9-10), 130-136.

Yap, S. F., Xu, Y., & Tan, L. (2021). Coping with crisis: The paradox of technology and consumer vulnerability. *International Journal of Consumer Studies*, 45(6), 1239-1257.