



Dario Krpan Dima Sayess Fatima Koaik George Philippe Farajalla Pujen Shrestha Robin Schnider March 4th, 2025

How agentic AI can be applied to behavioural science

Traditional AI models have been extensively studied for their influence on human behaviour. But agentic AI remains largely unexplored in behavioural science. Dario Krpan, Dima Sayess, Fatima Koaik, George Philippe Farajalla, Pujen Shrestha and Robin Schnider discuss the concept of agentic AI and explore the risks and opportunities of its responsible integration into behavioural science.

Artificial intelligence is transitioning from passive, instruction-based systems to autonomous entities capable of independently defining and pursuing goals. This agentic AI represents the next frontier in the technology. The innovation could revolutionise industries by enabling technology to make complex decisions and solve problems without direct human oversight.

Yet, despite its promise, agentic AI remains largely unexplored in behavioural science. Traditional AI models have been extensively studied for their implications in shaping human behaviour. For example, **digital footprints on social media** can be used to predict personality and other individual characteristics to inform personalised interventions. However, the impact of truly autonomous AI agents remains unknown.

Definition and current trends

Agentic AI marks a shift from the traditional version—which follows predefined instructions—to autonomous systems that independently set and pursue complex goals. Unlike conventional AI, which relies on human oversight and structured inputs, agentic AI makes decisions and optimises resources with minimal intervention. Its ability to **operate independently** in dynamic environments makes it valuable in fields requiring continuous problem-solving, such as healthcare and finance, while also contributing to industrial automation.

What distinguishes the agentic from the generative versions of AI is its focus on action rather than content creation. Generative AI, such as **ChatGPT** or **DALL·E**, is designed to produce text or images based on learned patterns, but it requires human prompts and lacks autonomy. In contrast, agentic AI is goal driven. It can form strategies and execute tasks over extended periods. While generative AI responds to input, agentic AI **actively seeks solutions**, making it more suited for complex, decision-intensive tasks. This distinction has led to an increasing focus on developing AI systems that do not just generate outputs but also independently take meaningful actions in various domains, from business to science.

Multi-agent systems

The rapid advancement of this technology is reshaping industries and raising critical policy considerations. One major development is the rise of **multi-agent systems**, where multiple AI entities work together to achieve shared objectives. These systems are already being deployed in logistics, where AI agents optimise supply chains, and in scientific research, where they independently generate and test hypotheses.

Another major trend is the improvement of long-term planning. With better reasoning capabilities, reinforcement learning and memory-based **architectures** allow agentic AI to refine its decisions over time, making it increasingly valuable in areas like investment management and crisis response, where rapid adaptation to shifting circumstances is essential.

Industries are beginning to **integrate agentic AI** into core operations. In finance, AI-driven trading systems autonomously respond to market fluctuations, optimising investments in real time. In healthcare, AI-powered patient monitoring systems detect potential health risks before they escalate. In robotics, agentic AI enables machines to operate independently in hazardous environments such as disaster zones or deep-sea exploration. These early use cases highlight the potential for AI systems that not only assist humans but take on complex responsibilities traditionally requiring human judgment. However, high-stakes applications raise concerns about data privacy and the need for strong regulatory safeguards.

While AI's ability to self-direct and adapt unlocks new possibilities across industries, its growing autonomy demands proactive governance. As these systems become more powerful and integrated into daily life, balancing innovation with ethical and legal boundaries will be key to ensuring that agentic AI serves human interests effectively.

Implications for behavioural science

AI's potential for transforming behavioural science is significant. Personalised interventions will allow agentic systems to automate behaviour change strategies and scale up research

methodologies. Traditional AI primarily supports structured decision-making. In contrast, agentic technology can autonomously identify behavioural patterns and adjust its approach based on user responses. This ability makes the technology valuable for public policy applications where continuous adaptation is required, as is the case in health interventions and financial decision-making.

However, as agentic AI becomes integrated into decision-making and intervention frameworks, policymakers must ensure that governance structures keep pace with the technology's capabilities, to prevent unintended consequences.

Automated behavioural interventions

Health

One of the most promising applications of agentic AI in behavioural science is personalised health coaching. Recent work has demonstrated how multi-agent AI systems can provide adaptive **nutrition and wellness coaching**, integrating behavioural science principles to tailor interventions to individual needs. For example, an AI-driven Barrier Identification Agent engages in motivational interviewing to uncover psychological and situational obstacles preventing users from adopting healthier eating habits. A second Strategy Execution Agent then provides customised behavioural change techniques, drawing from established frameworks such as Capability-Opportunity-Motivation-Behaviour (COM-B) and Behaviour Change Wheel (BCW) **models**.

By adjusting recommendations based on user feedback, AI agents enhance the effectiveness of interventions while reducing reliance on human coaches. This application showcases how agentic AI can improve accessibility and scalability in health and wellness programs. Ensuring data privacy and safeguards against unintended harms will be crucial for responsible implementation in health and wellness programs.

Financial decision-making

Another critical area where agentic AI could feed automated behavioural interventions to drive innovation is in financial decision-making. Behavioural science has long demonstrated that, due to cognitive biases, people struggle with issues such as saving for retirement or reducing debt. An agentic system could function as an autonomous financial advisor. For example, it could monitor user spending habits to predict risky behaviours. This could inform real-time nudges to encourage better financial choices.

Unlike static financial planning tools, an agentic system could continuously refine its approach based on user preferences, psychological barriers, and external economic conditions, making

financial decision-making more intuitive and tailored to individual needs. However, safeguards are needed to ensure consumer protection if AI-driven decisions result in financial losses or unintended risks.

Workplace and the future of employment

In the workplace, agentic AI could act as an adaptive coach, enhancing productivity and learning. AI-driven virtual mentors could monitor workplace behaviours to identify skills gaps and provide personalised feedback, thus improving performance. These systems could integrate behavioural science strategies such as temptation bundling (pairing an enjoyable task with a necessary but less appealing one) and mental rehearsal (helping employees visualise successful task execution) to increase motivation and engagement. By autonomously adapting to employees' learning styles and work habits, agentic AI could play a key role in organisational behaviour and performance management. However, safeguards are needed to prevent excessive surveillance and biased decision-making.

Behavioural science research

Beyond direct interventions, agentic AI could revolutionise behavioural science research itself. Traditionally, studying human behaviour requires extensive data collection and manual analysis. AI agents could autonomously generate and test behavioural interventions to refine psychological models based on real-time data. For example, AI-driven synthetic populations could be used to simulate human decision-making in high-stakes environments, allowing researchers to explore scenarios that would be difficult or unethical to test with real participants. This approach could accelerate discoveries in various areas of behavioural science, from health to sustainability.

Ethical and policy considerations

As agentic AI becomes increasingly integrated into behavioural science, ensuring transparency and accountability is essential. One of the most pressing concerns is the potential for AI to manipulate choices without informed consent. If AI autonomously adjusts behavioural strategies, it could reinforce cognitive biases or subtly nudge individuals toward decisions that do not align with their best interests. For example, an AI-powered financial advisor could prioritise investment options that maximise platform profits while framing them as the most suitable choices for the user, exploiting cognitive biases like loss aversion without explicit disclosure.

To mitigate these risks, robust regulatory frameworks must be developed to ensure AI operates within ethical boundaries. Key safeguards should include algorithmic audits to detect and correct biases in hiring and lending decisions, encryption protocols to protect sensitive user data from

cyber threats, and independent review boards to oversee AI-driven decisions in areas like healthcare and criminal justice. Maintaining public trust requires transparency in AI decision-making processes and clear accountability structures. A well-defined AI policy should balance innovation with ethical responsibility, ensuring that agentic AI enhances human well-being rather than undermining autonomy or deepening inequalities.

Final word

Agentic AI offers a powerful new tool for behavioural scientists, allowing for personalised and adaptive interventions. By applying behavioural science principles to AI-driven systems, we can enhance human decision-making in various domains, from health to education. However, realising this potential requires careful oversight to ensure these systems align with ethical standards and truly serve human well-being.

Sign up for our weekly newsletter [here](#).

- *This blog post represents the views of its author(s), not the position of LSE Business Review or the London School of Economics and Political Science.*
- *Featured [image](#) provided by Shutterstock.*
- *When you leave a comment, you're agreeing to our [Comment Policy](#).*

Author affiliations

Dario Krpan is an Assistant Professor in Behavioural Science in the Department of Psychological and Behavioural Science at LSE.

Dima Sayess is a Partner and Lead of the Ideation Center at Strategy& (part of the PwC network) Middle East.

Fatima Koaik is a Principal and Director of Behavioural Economics at the Ideation Center, Strategy& (part of the PwC network) Middle East, and a PhD alumna of the Department of Psychological and Behavioural Science at LSE.

George Philippe Farajalla is a Senior Associate and Data Scientist at the Ideation Center, Strategy& (part of the PwC network) Middle East.

Pujen Shrestha is a Behavioural Scientist and Data Scientist at the Ideation Center, Strategy& (part of the PwC network) Middle East.

Robin Schnider is a Senior Behavioural Economist at the Ideation Center, Strategy& (part of the PwC network) Middle East.

About the author

Dario Krpan

Dima Sayess

Fatima Koaik

George Philippe Farajalla

Pujen Shrestha

Robin Schnider

Posted In: LSE Authors | Management | Technology



© LSE 2025