

THE HUMAN TEST OF AI RESILIENCE

AI systems may work as designed and still weaken organisational judgement and trust. The real test of resilience lies in how people use, question and remain accountable for them.

- 1 *AI adoption is more than a technical implementation challenge. It changes how people exercise judgement, assign responsibility and understand their own expertise.*
- 2 *Resilient AI adoption requires leaders to anticipate how people will interpret a system before rollout, prepare them to exercise judgement with it and adapt governance as real patterns of use emerge.*
- 3 *As organisations learn to live and thrive with AI, adaptation depends on understanding how people actually behave—not how they ought to behave—and designing AI adoption around it.*

A regional bank is preparing to introduce an AI-assisted credit assessment tool for small-business lending. The system has been tested for months. It can process applications faster than human officers, flag inconsistencies in financial records and draw on alternative data to assess businesses with thin credit histories. For leaders, the business case seems obvious—faster decisions, lower processing costs and wider access to finance.

Then the tool meets the organisation.

Some loan officers worry that the model will make their judgement less valuable. Others quietly ignore its recommendations when they do not match their intuition. A few rely on the model too heavily, treating its input as a decision rather than an aid to decision-making. Customers question why a long-standing relationship manager can no longer explain a rejected application. Compliance teams want clearer documentation. Branch managers, under pressure to meet approval targets, begin to differ in how strictly they follow the model's guidance.

Nothing dramatic has happened; the system has not crashed, the model has not failed. Yet, the organisation is already becoming less coherent in how it uses AI.

This is the central challenge of AI resilience. The question is not whether an AI system is accurate or efficient. It is whether the organisation can absorb it into work without eroding trust, accountability, human agency or capability. Indeed, AI changes what tasks get done, but it also changes who gets heard, what counts as expertise, how people justify decisions and where responsibility seems to sit when outcomes are contested.

Though resilience is often discussed in relation to things like supply chains and geopolitical shocks, in the age of intelligent tools, it also carries an internal behavioural dimension. Organisations must adjust their approach to AI as it changes how people think and work.

A resilient AI strategy therefore rests on three disciplines: *anticipate* how people will interpret AI use before deployment; *prepare* employees to exercise judgement with AI; and *adapt* governance as real patterns of use emerge.

AI ADOPTION AS A BEHAVIOURAL CHANGE CHALLENGE

AI is appealing because it promises control amid uncertainty. It can detect patterns in volatile markets, optimise logistics across congested networks and help organisations respond more quickly to sudden changes in demand. The more complex the environment becomes, the more tempting it is to believe that better prediction will produce better control.

Yet many AI initiatives disappoint because they are built around a thin understanding of adoption. It is thought that people will accept a tool if it is presented as useful and shown to have real benefits. Human behaviour, however, rarely follows this script.

A doctor for instance does not evaluate a diagnostic model only by its accuracy. They ask what it means for clinical judgement or what it means for their liability. Likewise, a public officer using automated triage must still face the citizen seeking an explanation for how their case was routed. These are not minor implementation details, but are, in fact, the substance of adoption.

This is why AI adoption should be understood as a behavioural change management effort.¹ AI systems increasingly recommend, rank, interpret, allocate attention and sometimes act with a degree of autonomy. They enter the terrain of expertise. They often look authoritative even when their limits are poorly understood. They can also make employees feel that the knowledge and discretion that once defined their work are being displaced.

A behavioural approach to human-centred AI begins from the premise that people are not perfectly rational users.² They bring habits, fears, pride, biases, professional identities and social concerns to every interaction with intelligent tools. Even with full

transparency and a clear business case, adoption can still falter if the human factor is not taken into account. Designing for the predictable messiness of human behaviour therefore forms the foundation of resilient AI adoption.

ANTICIPATION: PERCEPTIONS BEFORE THE ROLLOUT


Before an AI system goes live, leaders usually know a great deal about the technology: they know the vendor, implementation cost, integration plan, data requirements and expected (or rather, supposed) efficiency gains. They may, however, know far less about the social meaning of the system they are about to introduce.

That meaning matters because an AI tool described as “decision support” might be seen as a signal that management no longer trusts employees’ judgement. A performance dashboard to promote sales may be experienced as surveillance instead of motivation. A recommendation engine could be intended to reduce errors, but employees quietly wonder whether they will be blamed for following the system, or overriding it.

These interpretations are early evidence of where the organisation is vulnerable. AI can improve some elements of performance while also unsettling the human conditions that make broad performance sustainable. When employees experience AI as replacing rather than supporting judgement, adoption becomes fragile. When customers, clients or citizens encounter decisions shaped by AI, they may still expect reasons that a human can explain and defend. The risk, then, is not that people reject AI, but that employees – and by extension the organisation – are unprepared for how AI reshapes roles and expectations.

The practical exercise here is not stakeholder engagement in the generic sense – it is behavioural anticipation. Leaders should identify the moments where AI is likely to disturb the status quo. A model that forecasts inventory needs raises different concerns from one that screens job applicants, assesses loan eligibility or allocates healthcare resources. The more directly a system affects people’s opportunities or dignity, the more important it becomes to anticipate not only how accurate it is, but how people will respond to its use.³

A useful test is to examine how a serious disagreement between an experienced employee and



the AI system would unfold. The employee believes the recommendation is wrong. The model appears confident. A manager is watching. A customer is waiting. What happens next? If the employee cannot understand the basis of the recommendation, challenge it safely or escalate the case without appearing obstructive, the organisation has not designed for AI resilience; it has designed for compliance on the surface and confusion underneath. Anticipation is the discipline of discovering these weak points before they become habits of avoidance or over-reliance. This, in effect, is what resilience means here: the organisation’s capacity to absorb friction between human judgement and AI output without fractures emerging.

PREPARATION: RECOGNISE LIMITATIONS, REHEARSE JUDGEMENT

Preparation is often reduced to training – learning what to click and type on a given checklist of different scenarios. This may be enough for ordinary software, but it is rarely enough for AI.

When AI enters judgement-heavy work, people need more than instructions – they need practice in deciding how much weight to give the system, when to question it and how to explain decisions shaped by it.

This is challenging because people can go wrong in opposite directions. Some become sceptical after seeing an algorithm make a mistake, even when it performs better than humans on average.⁴ Others defer too readily because the output appears quantified, is produced almost instantaneously and looks sophisticated or difficult to challenge.⁵ In a busy workplace, a checklist approach for what to do can become the path of least resistance – a quick, convenient answer to a problem that employees have neither time nor bandwidth to address.

Good preparation surfaces these tendencies. Consider a hospital introducing an AI imaging tool. A thin version of preparation would announce that the system has been validated and provide a short training session. A stronger version would show clinicians where the system performs well, where it struggles, how uncertainty is displayed and when human review is mandatory. This would allow clinicians to compare their judgement against the model’s output in practice. It also gives them a sense of when they need to be *even more* vigilant. In banks, relationship managers would be able to explain AI-supported credit decisions without hiding behind the model. In logistics, planners can understand where local knowledge should override AI-powered recommendations. Interfaces should

also be treated as behavioural environments—how an interface is designed can affect how users make decisions. Default options, user experience, presentation of social norms can all steer users into a particular direction. The key question is whether interfaces are equipped to help employees engage meaningfully, beyond simply being kept in the loop. Preparation therefore means building the conditions for judgement.

ADAPTATION: OBSERVE, LEARN AND ADJUST

The launch of an AI system is often treated as the end of implementation. In reality, it is the start of discovery. Only after launch does the organisation learn how people actually use the system, especially when targets, time pressure and customer demands are all in play.

People adapt around AI. They create shortcuts, discover exceptions, develop trust or distrust. They find ways to satisfy the system’s requirements while preserving old routines (such as by performing boundary work). They may use AI for purposes beyond its original design, or they may become dependent on it in ways that weaken underlying capability.

This is where governance often falls short. Many organisations have approval processes, validation checks and policy documents. Though necessary, they typically capture the system at a point in time. Behaviour is constantly evolving. As an example, a productivity tool might gradually change expectations about

workload and availability, and so workers might evolve their use patterns to match growing expectations about what they are supposed to produce.

Resilient organisations treat actual use as evidence. By observing the various ways employees engage, or avoid, AI systems, organisations can discover whether and how AI unintentionally changes the meaning of good work.

Employees often pick up on this change in meaning before formal governance systems do. A nurse could notice that an AI-supported output confuses elderly patients. A claims officer may see that a fraud detection tool flags certain customer profiles too often. A procurement manager might discover that generative AI produces supplier summaries that sound convincing but leave out crucial disclaimers. These observations are key early-warning signals.⁶

Adaptation also requires accountability to be revisited as practice evolves. When decisions are shaped by both humans and AI, responsibility can blur. If a manager follows an AI recommendation that later proves harmful, who is accountable? Is it the model developer, the vendor, the user, the department head or the organisation? A resilient governance system does not wait for a failure to answer these questions. It clarifies all the necessary rights, escalation routes and documentation standards that need to be considered beforehand, then revises them as new cases expose new ambiguities.

This is the difference between AI governance as compliance and AI governance as organisational learning. Compliance asks whether the approved process was followed, whereas learning asks what the organisation is discovering about the interaction between people, tools

and consequences. In fast-changing environments, that second question is the one that builds resilience.

LEADING RESILIENT AI ADOPTION

Effective adoption of AI begins with approaching AI as change management rather than technology implementation. That shift changes the work leaders must do.

First, leaders should identify the behavioural stakes of AI adoption. A significant rollout of an AI system should come with an account of the expected behavioural changes. Leaders need to clearly communicate these expected changes to employees from the outset.

Second, leaders should embed behavioural risk mapping within the deployment process. Before a system goes live, the organisation should examine where distrust, over-trust, anxiety, workarounds or accountability gaps are likely to arise. This should be tailored to the specific technology in question—a tool for drafting marketing copy, for example, creates very different behavioural risks from a tool for performance evaluation or credit scoring. Mapping these risks further strengthens leaders’ communications, as they can shape expectations around likely scenarios and concerns before they emerge.

Third, user involvement should happen earlier than it often does. Too many employees are invited to test systems only after key design choices have been made. Genuine involvement means understanding how workflows could change before final design changes have been made. Involving employees at

the design stage allows them to meaningfully participate in the change initiative while providing direct feedback from end-users. Their involvement is thus both symbolic (as employees feel heard) and functional.

Lastly, governance must stay close to practice. Committees, audits and policy documents are useful only if they remain connected to what people are doing on the ground. Leaders should regularly ask where the system is being bypassed, where it is being followed and where it is creating new pressures. The answers may be uncomfortable. They are also the raw material for resilience.

FROM ROLLOUT TO RESILIENCE

The promise of AI is real. It can help organisations respond to volatility, expand access, improve decisions and operate at a scale that would otherwise be impossible.⁷ Yet, poorly managed AI can also make organisations brittle. It can centralise decisions without clarifying accountability, accelerate work without preserving judgement or create an appearance of objectivity while hiding highly contested assumptions. In short, it can make organisations faster, but less trusted and stable.

The organisations that thrive with AI will be those that understand the paradox at the heart of intelligent technology: the more powerful AI becomes, the more important human behaviour becomes. Leaders must design not only for what machines can do—and the efficiency they bring—but for how people will respond to them.

Resilient AI starts with anticipation, accounting for the human meaning of a system before it is introduced. It relies on preparation, where leaders rehearse judgement rather than simply train usage. And it requires adaptation, as organisations treat real-world use as evidence for how governance must evolve.

In uncertain times, organisations often look for resilience in stronger systems, better forecasts and faster tools. These matter. But the most resilient AI systems may ultimately be those designed not around perfect machines, but around perfectly irrational and deeply human users. 



JACK MCGUIRE is an Assistant Professor of Organisational Behaviour and Human Resources at Singapore Management University. His research examines how emerging technologies, leadership, identity and emotions shape the way people work. His current work explores topics including the intersection between artificial intelligence and leadership, awe in organisations and how workers imagine their role in the future of work.

For a list of references to this article, please visit <https://smu.sg/SMUAMIAug2026> or scan the QR code below.



If an employee cannot understand nor challenge AI outputs safely, the organisation has not designed for AI resilience, but for compliance.