

# AI Does Eliminate Work. It Leaves Us With the Decisions.

Why the biggest risk in AI implementation is not the technology, but the quality of the human decisions that govern it.

Sylvain Rottier

Co-Founder, Mindmarqs | Psychologist

July 2026

---

## Abstract

As AI automates execution, human work increasingly shifts toward reviewing, approving and governing AI-generated decisions. This paper introduces the concept of decision density and explains why successful AI transformation depends on an organization's cognitive readiness, not just its technical readiness.

It outlines why measuring decision-making capability is becoming essential for AI governance, leadership development and organizational performance. As routine execution is delegated to intelligent systems, the remaining work increasingly consists of evaluating recommendations, resolving ambiguity, governing exceptions and accepting accountability for decisions executed at machine speed.

The organizations most likely to realise sustained value from AI will not necessarily be those deploying the most advanced technology. They will be those that understand how their people make decisions, where cognitive vulnerabilities exist and how those capabilities can be strengthened over time.

## Ideal for

AI Transformation Leaders, HR Executives, Talent Management, Leadership Development, Digital Transformation, Executive Assessment.

## You'll learn

- Why AI dramatically increases the number and importance of human judgement calls.

- What “decision density” means for leadership and workforce capability.
- Why measuring cognitive readiness should become part of every AI implementation programme.

## Key Insights

- **The Decision Density Shift:** AI removes execution work while dramatically increasing the number, speed and consequence of human judgement calls across every function.
- **The Governance Challenge:** many AI implementation failures stem not from technology, but from unclear accountability and insufficient human oversight.
- **The Automation Bias:** humans consistently over-trust confident AI recommendations, increasing the risk of poor decisions when independent judgement is not actively maintained.
- **The Hidden Measurement Gap:** organizations assess skills, experience and personality, but rarely evaluate how people actually make decisions under pressure.
- **The Cognitive Readiness Imperative:** successful AI adoption depends as much on attention, cognitive resilience and judgement as it does on technology and process design.
- **The Assessment Opportunity:** behavioural measurement enables organizations to identify, develop and strengthen decision capability before poor judgement is amplified by AI.

## Section 1 – AI isn't removing decisions. It's multiplying them.

The most significant insight behind this research did not emerge solely from the academic literature. It developed through a recurring pattern observed during discussions with senior leaders implementing agentic AI across a broad range of industries. Although the operational contexts differed considerably, the underlying challenge proved remarkably consistent.

Most AI implementation programmes began with the expectation that automation would reduce human workload by removing routine execution while leaving only occasional exceptions requiring human intervention. In practice, organizations are experiencing a different outcome. Rather than reducing the importance of human judgement, AI is concentrating on it. As routine execution is delegated to intelligent systems, the remaining work increasingly consists of evaluating recommendations, resolving ambiguity, governing exceptions and accepting accountability for decisions executed at machine speed.

This shift is evident across a wide range of knowledge-intensive roles. Marketing professionals who previously developed campaign strategies now evaluate dozens of AI-generated alternatives before determining which deserve investment. Insurance claims specialists increasingly review automated recommendations rather than manually building each case, while logistics coordinators supervise AI-driven optimisation platforms that continuously generate routing decisions requiring periodic human intervention. Although the operational responsibilities differ, each role reflects the same structural transition: execution is progressively automated while human judgement becomes more frequent, more consequential and more strategically significant.

This emerging pattern can be described as **decision density**: the number, pace and consequence of judgement calls required within a role over a given period of time. As organizations deploy increasingly capable AI systems, decision density rises naturally because every task delegated to AI leaves behind fewer operational activities but proportionally more approvals, exceptions, escalations and governance decisions. Human judgement increasingly becomes the control point through which AI execution flows.

The implications extend well beyond productivity. Historically, poor decisions were moderated by slower execution, multiple layers of review and operational constraints. Agentic AI substantially reduces these buffers. Poor judgement can now be propagated rapidly and consistently across an entire workflow, just as high-quality judgement can generate value at unprecedented scale. Organizational performance therefore becomes increasingly dependent on the quality of the human decisions that govern intelligent systems.

This represents a fundamental change in the cognitive demands of work. Many roles that were historically execution-oriented now require sustained attention, critical evaluation, contextual reasoning and the confidence to challenge highly plausible AI-generated recommendations. Yet while organizations have invested heavily in AI platforms, governance frameworks and process redesign, relatively little attention has been given to whether the workforce possesses the cognitive capabilities required to perform these emerging responsibilities effectively.

The question is no longer whether organizations are technically prepared to deploy AI. Increasingly, the defining challenge is whether their people are cognitively prepared to govern it.

## Section 2 – When decision density becomes reality

The implications of increasing decision density became particularly clear to me long before generative AI entered mainstream business. They emerged during the introduction of autonomous mobile cleaning robots across Europe, the Middle East and Africa. At the time, the focus was understandably on the technology, its capabilities, deployment and operational performance. Far less attention was given to the people whose roles would fundamentally change as a consequence.

Before automation, the supervisor's responsibilities were relatively straightforward. Success was largely determined by coordinating people, scheduling work and confirming that cleaning activities had been completed to the required standard. Decision-making was largely local, immediate and operational.

The introduction of autonomous robots transformed that role in ways few anticipated.

Every robot generated continuous operational data. Cleaning schedules became dynamic rather than fixed. Exceptions became visible in real time. Individual decisions regarding scheduling, prioritisation and intervention no longer affected a single task; they influenced an interconnected system operating across multiple locations and stakeholders. Activities that had previously relied on local judgement evolved into decisions with broader operational, financial and customer implications.

Managing a fleet of intelligent machines therefore required a fundamentally different cognitive capability from managing the people performing the work manually. The role shifted from coordinating activities to governing a system. Success depended less on operational supervision and increasingly on interpreting information, recognising patterns, anticipating consequences and making sound decisions within a far more complex environment.

One lesson became particularly clear. A poor decision was no longer contained. When an incorrect judgement was embedded within an automated workflow, the system repeated that decision consistently and at scale until someone recognised the problem and intervened. Automation did not simply increase efficiency; it amplified the consequences of both good and poor judgement.

Looking back, the technology itself was not the greatest implementation challenge. The real challenge was that the cognitive demands of the supervisory role had changed almost overnight, while the preparation provided to those assuming these new responsibilities had barely changed at all. Considerable investment had been made in deploying the robots. Very little had been made in preparing people to govern them.

What was observed then is now unfolding across almost every knowledge-intensive function adopting agentic AI.

The experience of Klarna illustrates this shift at a much larger scale. In 2024, the company announced that its AI assistant was performing work equivalent to

approximately 700 customer service agents and handling more than two million customer conversations within its first month. Yet within a year the organization had begun increasing human recruitment once again. The challenge was not the AI's ability to manage volume. It was the complexity of the decisions that volume had previously concealed, when to escalate an issue, when to exercise discretion, how to respond when a customer situation fell outside predefined rules, and ultimately what standards of judgement represented the organisation's values.

These decisions had always existed. Previously they were embedded within the execution of everyday work. As execution became increasingly automated, they emerged as the primary responsibility of the remaining human roles.

This pattern is now reflected in broader industry research. Gartner predicts that more than 40 per cent of agentic AI initiatives will be cancelled before the end of 2027, not because the technology fails to perform, but because organizations struggle to establish effective governance, accountability and decision oversight (Gartner, 2025).

The emerging constraint is therefore no longer technological capability. It is the capacity of people to make consistently sound decisions in environments where those decisions are executed immediately, repeatedly and at machine scale.

### Section 3 – The missing measure of AI readiness

Organizations invest considerable effort in assessing their people. Recruitment processes evaluate technical expertise, experience, personality, leadership potential, cognitive aptitude and cultural fit. Leadership programmes assess competencies, emotional intelligence and behavioural styles. Yet one capability remains remarkably absent from most assessment frameworks: how people actually make decisions.

Most assessment approaches rely heavily on self-report. Candidates describe how they believe they behave, leaders explain how they think they make decisions, and organizations infer future performance from questionnaires, interviews and psychometric profiles. While these methods provide valuable insights into preferences and behavioural tendencies, they reveal relatively little about the cognitive processes that determine judgement under real conditions of uncertainty, complexity and time pressure.

This limitation has been recognised for decades. Kahneman and Tversky's pioneering work on heuristics and biases demonstrated that much of human judgement is shaped by cognitive processes operating below conscious awareness (Kahneman & Tversky, 1974). People are often unable to accurately explain how they reached a decision because many of the mechanisms influencing attention, information processing and judgement occur automatically. Self-report, therefore, cannot reliably capture the quality of decision-making itself.

The emergence of agentic AI has transformed this long-standing limitation into an urgent organizational challenge.

A growing body of research demonstrates that human-AI collaboration does not automatically improve decision quality. Vaccaro and colleagues' (2024) meta-analysis of 106 experimental studies found that human-AI teams frequently performed worse than the stronger decision-maker operating alone, particularly on judgement-intensive tasks. One of the principal causes is **automation bias**, the well-documented tendency for people to place excessive trust in confident, authoritative technological systems, even when independent evaluation would produce better outcomes (Parasuraman & Manzey, 2010).

The implications extend beyond isolated decision errors. Recent workforce research suggests that fewer than one in eight employees currently possess the combined capabilities required to effectively oversee AI-generated outputs in complex decision environments (Global Data Literacy Benchmark, 2025). As organizations increase their reliance on autonomous systems, the quality of human oversight becomes an increasingly significant determinant of organizational performance.

Our own behavioural assessment data points to a similar conclusion. The greatest differences between effective and ineffective decision-makers are rarely explained by domain knowledge or technical expertise alone. Instead, they are associated with cognitive capabilities such as sustained attention, information search behaviour, resilience under increasing cognitive load and the willingness to challenge plausible but potentially incorrect recommendations.

One pattern has proven particularly consistent. Individuals making the highest-quality decisions rarely accept the first credible answer presented to them. Instead, they deliberately search beyond their initial informational comfort zone, seek evidence that challenges rather than confirms their assumptions and maintain independent evaluation even when an apparently authoritative recommendation is readily available. Others, particularly under increasing information volume or time pressure, display what can be described as **cognitive surrender**, accepting AI-generated recommendations not because they have been critically evaluated, but because sustained independent judgement becomes cognitively demanding.

These behavioural differences are difficult to detect through conventional assessment methods because they emerge during the decision process itself rather than through what individuals subsequently report about that process.

The organizational implications extend beyond individual performance. Every organization develops a characteristic 'decision culture', a collective pattern of attention, judgement and behavioural heuristics that shapes how decisions are made throughout the enterprise. Research by Hodgkinson and Healey (2023) demonstrates that these decision patterns remain remarkably stable over time, while Barsade and colleagues (2018) showed that behavioural patterns spread throughout organizations

primarily through leadership behaviour rather than formal policy.

Historically, the consequences of these invisible decision cultures were moderated by slower execution, multiple review stages and distributed operational responsibility. Agentic AI removes many of these natural constraints. As execution becomes increasingly automated, organizational outcomes become progressively more dependent upon the quality of the decisions that guide those systems.

The challenge is therefore no longer simply understanding how people think. It is developing reliable ways to observe how they actually decide.

## Section 4 – Cognitive readiness: the missing dimension of AI transformation

Throughout discussions with organizations implementing agentic AI, one observation has remained remarkably consistent. Considerable attention is devoted to technology selection, system integration, governance frameworks and change management. These are all essential components of successful AI transformation. Yet one fundamental question is rarely addressed.

Are the people responsible for governing AI cognitively prepared for the decisions it leaves behind?

Traditional implementation programmes typically assess technical readiness, process maturity and organisational capability. They evaluate whether the technology functions as intended and whether appropriate governance structures are in place. Far less attention is given to the individuals who remain accountable for the quality of AI-assisted decisions. Yet as autonomous systems assume an increasing share of execution, those individuals become the principal determinant of organisational performance.

Cognitive readiness describes an organisation's capacity to sustain high-quality judgement in environments characterised by increasing decision density. It encompasses the behavioural capabilities that enable individuals to maintain attention, evaluate evidence critically, recognise uncertainty, challenge plausible recommendations and remain resilient under sustained cognitive load. These capabilities are becoming as important to successful AI adoption as technical competence or digital literacy.

This changes the nature of workforce readiness. Historically, organisations prepared employees to perform tasks. Increasingly, they must prepare them to govern intelligent systems. The distinction is significant. Technical skills enable people to operate AI. Cognitive capabilities determine whether AI is applied wisely.

This raises a different set of organisational questions.

Which business functions are experiencing the greatest increase in decision density? Where are employees expected to review hundreds of AI-generated recommendations rather than produce the work themselves? Which roles demand sustained attention and independent judgement for prolonged periods? Where is automation bias most likely to emerge? Which teams demonstrate the cognitive resilience required to maintain decision quality as AI becomes more deeply embedded within operational workflows?

These questions cannot be answered through traditional competency models or conventional psychometric assessments alone because they concern the quality of decision-making rather than knowledge, personality or experience. They require organisations to observe how decisions are actually made under realistic conditions, identifying behavioural patterns that influence judgement before they become embedded within AI-enabled workflows.

This is where behavioural measurement offers a significant opportunity. Rather than asking individuals how they believe they make decisions, behavioural assessment enables organisations to observe attention, information search, evaluation strategies and decision processes directly. The objective is not simply to identify stronger decision-makers, but to understand the cognitive strengths and vulnerabilities that exist across teams, functions and the organisation as a whole.

Mindmarqs was developed around this principle. Its behavioural assessment methodology was originally designed to identify individuals capable of making consistently high-quality decisions in complex, high-stakes environments. The emergence of agentic AI has significantly broadened its relevance. The same behavioural capabilities that distinguish exceptional decision-makers are now becoming increasingly important across a much wider range of organisational roles, as AI shifts human work from execution towards judgement and governance.

The organisations most likely to realise sustained value from AI will not necessarily be those deploying the most advanced technology. They will be those that understand how their people make decisions, where cognitive vulnerabilities exist and how those capabilities can be strengthened over time. As AI continues to evolve, cognitive readiness will increasingly become a defining component of organisational readiness itself.

## Conclusion

Artificial intelligence is fundamentally changing the relationship between people and work. Routine execution is increasingly delegated to intelligent systems, while the remaining human responsibilities become progressively more centred on judgement, oversight and accountability. As this transition accelerates, the success of AI initiatives will depend less on the capability of the technology itself and increasingly on the capability of the people responsible for governing it.

This represents a significant shift in how organisations should think about workforce readiness. Technical readiness remains essential, but it is no longer sufficient. The ability to sustain attention, evaluate information critically, challenge confident recommendations and exercise sound judgement under increasing decision density is rapidly becoming a strategic organisational capability.

For many organisations, this capability remains largely invisible. Existing assessment frameworks evaluate knowledge, skills and behavioural preferences, yet provide limited insight into the cognitive processes that ultimately determine decision quality. As AI removes the execution work that previously absorbed much of organisational activity, these underlying patterns become more exposed, and their impact more significant.

The next generation of AI transformation will therefore be defined by more than technology adoption. It will be shaped by an organisation's ability to understand, measure and develop the human judgement that guides intelligent systems. Technology may increasingly determine what organisations *can* do. Cognitive capability will determine how well they do it.

The future of work will not be limited by artificial intelligence.

It will be defined by the quality of the human decisions that direct it.

## References

- Barsade, S.G., Coutifaris, C.G.V. & Pillemer, J. (2018). Emotional contagion in organizational life. *Research in Organizational Behavior*, 38, 137–151.
- Brynjolfsson, E. (2022). The Turing Trap: The promise & peril of human-like artificial intelligence. *Daedalus*, 151(2), 272–287.
- Gartner. (2025). Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027.
- Global Data Literacy Benchmark. (2025). AI is Outpacing Human Competency.
- Hodgkinson, G.P. & Healey, M.P. (2023). The heuristics and biases of top managers. *Journal of Management Studies*, 60(6), 1593–1627.
- Kahneman, D. & Tversky, A. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Mollick, E. (2024). *Co-Intelligence: Living and Working with AI*. Portfolio/Penguin.
- Parasuraman, R. & Manzey, D. (2010). Complacency and bias in human use of automation. *Human Factors*, 52(3), 381–410.
- Vaccaro, A. et al. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8, 2293–2303.