International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 6, Nov-Dec-2026, PP: 01-5

Leading with AI: A Human-Centred Framework for Bridging Managerial Decision-Making and Artificial Intelligence

Ms Srija R, Assistant Professor
Prof. (Dr.) Sreekanth Sreedharan
Dr. M. Dhanalakshmi, Associate Professor
Department of Management Studies
LEAD College (Autonomous), Dhoni, Palakkad (Affiliated to the University of Calicut)

***Abstract* – Managers now take many decisions with an AI system close at hand, but whether those decisions are any better for it is far from settled. Some field studies report striking gains. Yet when 106 experiments were pooled, human-AI pairs did worse on average than the stronger of the two working alone, and the shortfall was clearest in decision tasks. We argue that the explanation is organisational: results depend on who is asked to do what. Drawing on work on bounded rationality, automation and human-AI collaboration, we classify managerial decisions using two questions. How structured and data-rich is the decision? And how costly and how permanent would a mistake be? The answers yield four ways of involving AI, which we call Automate, Supervise, Explore and Deliberate. Each works only when three conditions hold: a named person answers for the decision, trust is matched to the system's track record, and outcomes are routinely checked. We set out three testable propositions and consider what the argument means for Indian organisations under the 2025 governance guidelines. For leaders, we suggest, the real question is seldom whether to use AI. It is where a human being must still have the final word.**

***Keywords* - artificial intelligence, managerial decision making, human-AI collaboration, leadership, AI governance, bounded rationality**

I. INTRODUCTION

Ask any manager what the job consists of and sooner or later the answer comes round to decisions. Simon saw this long ago, and he also saw why the decisions are rarely ideal: a manager has only so much information, so many hours and so much attention, and therefore stops at an option that will do [1].

Computers relieved the first of these limits decades ago by putting more data on the manager's desk. AI is different because it relieves the other two. A model can work through more material in a few seconds than a person could in a year.

Agrawal, Gans and Goldfarb describe the change in economic terms: prediction has become cheap [2]. And what is cheap gets used. Banks score loan applicants, retailers forecast demand, hospitals rank patients for attention, HR teams screen CVs, and since 2023 nearly every office worker has had a text generator within reach. Behind all this spending lies an assumption that seems almost too obvious to question – that a machine which predicts well, working with a person who knows the business, must outperform either one alone.

Often it does not. Pooling 106 experiments, Vaccaro, Almaatouq and Malone found that human-AI pairs fell short of the stronger partner on average, and most clearly when the task was to make a decision [3]. Other studies, reviewed below, report large gains in real workplaces. In our reading, neither camp is wrong. The difference lies in whether the partnership was arranged well or badly.

If that is right, the question belongs on the leader's desk, not only with the IT department. Someone decides which decisions are handed to software, which are shared and which are kept. Someone signs for the result. Our aim is practical: to give that person a rule of thumb they can defend when challenged.

The paper makes three contributions. It brings together strands of research in management, psychology and economics that rarely cite one another. It proposes four modes of AI involvement, chosen according to how structured a decision is and how much an error would cost. And it closes with testable propositions and practical guidance written with Indian organisations in mind.

Section II reviews prior work and Section III describes our method. Section IV presents the framework, and Sections V and VI discuss its implications and conclude.

II. LITERATURE REVIEW

1. What Goes Wrong with Unaided Judgement

Bounded rationality [1] explains only part of the problem. The other part, which drew attention much later, is inconsistency. Give the same file to two underwriters – or to one underwriter before and after lunch – and the answers may differ. Kahneman, Sibony and Sunstein call this variability noise [4]. Software has many weaknesses,

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 6, Nov-Dec-2026, PP: 01-5

but noise is not one of them: the same inputs always yield the same output.

How much is consistency worth? Studying bail hearings in the United States, Kleinberg and his co-authors estimated that a model ranking defendants by risk could cut crime without detaining more people, or detain fewer people without raising crime [5]. Bail suits a machine for three reasons. The decision is made thousands of times, every case has the same structure, and the records go back years.

2. Who is Good at What

Jarrahi offers a distinction we find useful. When the difficulty is too much information, AI has the edge; when it is too little information, or disagreement about goals, people do better [6]. Starting from a similar point, Shrestha, Ben-Menahem and von Krogh describe three arrangements: delegate the decision fully to AI, let human and machine work in sequence, or combine their independent judgements [7].

Raisch and Krakowski add an important caution. The line between augmenting a person and replacing one does not stay where it is drawn. A task that is assisted this year may be fully automated next year, and an automated process may need people again when conditions change [8]. Any division of labour a leader settles on will have to be revisited.

3. Do People Listen to the Machine?

Good advice is useless if ignored, and bad advice is harmful if swallowed whole. Dietvorst, Simmons and Massey watched participants abandon an algorithm after seeing it slip, although it still beat the human forecaster they preferred [9]. Logg, Minson and Moore found something that looks like the reverse.

Ordinary participants leaned more heavily on a recommendation once they were told an algorithm had produced it; experts were less impressed [10].

We read these two studies as describing the two failure modes that Parasuraman and Riley identified in 1997: misuse, where people rely on automation too much, and disuse, where they ignore a system that works [11]. Both can appear in the same office, sometimes on the same afternoon. What organisations need is trust that follows the evidence – higher for the kinds of decision the system has handled well, lower for the rest.

4. What Field Studies Show

Laboratory results only go so far, so the recent workplace studies matter. In a customer-support setting with 5,172 agents, Brynjolfsson, Li and Raymond recorded about 14 per cent more issues resolved per hour once a generative assistant was introduced. Newer and weaker agents improved the most; veterans hardly changed [12]. Dell'Acqua and colleagues ran an experiment with 758 management consultants. On assignments that suited the tool, the AI group got through roughly an eighth more tasks, finished about a quarter sooner, and their output was graded over 40 per cent better. Then came a task deliberately picked to defeat the tool. There the AI group reached the right answer less often than the control group, by 19 percentage points [13].

The meta-analysis cited earlier, covering 370 effect sizes, is consistent with this pattern. Human-AI pairs helped when content was being created and hurt when a choice had to be made. Relative strength mattered too: where the human alone outperformed the AI, the pair improved on it, but where the AI alone was better, adding a human made results worse [3].

5. What is Missing

The literature, then, explains what machines do well, where people are still needed and why trust goes wrong. It does not tell a branch manager or a hospital administrator what to do about the decision in front of them on Monday morning. Existing typologies classify organisational designs [7] or types of task [6]. Neither gives much weight to the question managers actually worry about: what happens if this goes wrong? Section IV attempts to bring that question into the analysis.

III. METHODOLOGY

We should be clear about what kind of paper this is. It reports no survey or experiment of our own. The framework comes from reading and comparing published studies – an integrative review – an approach suited to a topic whose evidence is spread across several disciplines.

Sources were chosen purposively rather than through a systematic search. To be included, a work had to appear in a peer-reviewed journal or come from an established publisher or a government or standards body; it had to address decisions people make with the help of AI; and it had to be either widely cited or based on recent large-sample evidence. While reading, we kept two lists.

The first recorded every feature of a decision that an author linked to good or poor results from human-AI work.

The second recorded the safeguards authors recommended. Two features appeared again and again in the first list – how structured the decision was and how costly an error would be – and these became the basis of the framework.

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
ISSN:3048-7722
Volume 4, Issue 6, Nov-Dec-2026, PP: 01-5

The limitation is obvious. Different readers might have selected different sources, and nothing here has yet been tested against field data. We therefore present propositions, not findings.

IV. PROPOSED FRAMEWORK

1. Two Questions About Any Decision

We begin from the view that a leader can ask two questions about any decision without needing a data scientist in the room.

How structured is it? Some decisions come round daily, follow settled rules and leave a long trail of earlier cases whose outcomes are known. Sanctioning a small personal loan is like this; so is sending a customer complaint to the right desk. Others are one-off. The aims may be contested and the past is a poor guide. Whether to open in a new state, or how to respond when a plant has an accident, are decisions of that sort. Readers will recognise Jarrahi's contrast between complexity and uncertainty here [6].

What does a mistake cost, and can it be undone? If an online shop recommends the wrong kurta, the customer simply scrolls on and nothing is lost. If a qualified candidate is rejected, a patient is sent home too early, or a farmer is wrongly refused a loan just before sowing, someone is harmed, and the harm is hard to reverse.

2. Four Modes

Combining the two questions gives four modes, set out in Table I.

Table 1. Four modes of involving AI in a managerial decision

Mode	How structured?	Cost of a mistake	What the AI does	What the person does	Examples
Automate	Highly	Small, easily corrected	Takes the decision and acts on it	Writes the rules, checks a sample	Matching invoices, reordering stock, routing queries
Supervise	Highly	Large, hard to correct	Recommend and gives reasons	Examines and signs off each case	Loan sanction, fraud alerts, hospitaltriage
Explore	Loosely	Small, easily corrected	Produces options and rough drafts	Picks, tries out, improves	Campaign ideas, product concept first drafts
Deliberate	Loosely	Large, hard to correct	Provides analysis and argues the other side	Decides and answers for it	Strategy, senior appointments, crisis response

Automate. Software is more consistent than a clerk and cheaper, and an occasional error does little damage. Asking a person to approve every item wastes that person’s time. It is better to have them write the rules and check a sample each week.

Supervise. The model will probably still predict well, but the stakes are too high to let it act alone. Our concern here is the rubber stamp. If pairs do worse when the AI is the stronger partner [3], a reviewer adds value only by knowing something the system does not – what the borrower said in person, say, or how the local market has shifted. Reviewers also need to see why the system recommends what it does. A bare score invites a bare signature.

Explore. Data are thin, but so are the penalties, so there is little reason for caution. Let the tool generate twenty options and discard eighteen. The gains reported for creative and drafting work fall in this quadrant [3], [13].

Deliberate. Neither question favours the machine. These decisions are the most likely to fall outside what Dell'Acqua and colleagues call the frontier, where the tool’s answer is fluent, confident and wrong [13]. AI still has a role, but as a sparring partner: ask it to attack the preferred option, or to list what would have to be true for the plan to succeed. The decision itself stays with people who can be named.

3. Three Conditions That Hold the Modes Together

Sorting decisions into categories achieves little on its own. In our view, three conditions must hold whichever mode applies.

- Someone answers for it. Every AI-supported decision has a named human owner. Responsibility cannot be shifted to the software or to the vendor that supplied it.
- Trust follows the record. Staff are told how the system has performed on this type of decision and are trained to recognise both failure modes described earlier – over-reliance and neglect [11].

- Results are checked. What the system recommended, what the person decided and what eventually happened are recorded and compared. When the evidence changes, the decision moves to a different mode, as one would expect if the boundary of automation keeps shifting [8].

4. Propositions

Three propositions follow, each open to testing.

- P1. Organisations that match the mode to the structure and stakes of each decision will decide better than organisations that use one mode for everything.
- P2. Under the Supervise mode, human review improves outcomes only when the reviewer has relevant information the AI system does not have.
- P3. AI will improve decision quality more in organisations that track outcomes and move decisions between modes than in organisations that do neither.

V. DISCUSSION

1. What Leaders Should Take from This

Our first recommendation is unglamorous: before buying anything, list the decisions. A management team that writes down its recurring decisions and assigns each to one of the four modes usually finds it has been treating them all the same way – typically by making someone approve everything.

Skills raise a harder issue. Novices gain most from AI assistance [12], and no one would want to deny them that. But we see a cost that will surface later. Judgement is learned by using it, and a trainee who is always handed a draft answer may never develop it. In five years' time, who will be ready to take the Deliberate decisions? We believe leaders should deliberately keep some work unassisted for junior staff, much as a good teacher still sets sums to be done without a calculator.

Finally, the leader's own role changes. It involves less deciding and more designing how decisions are made: which data a system may use, what a reviewer is expected to contribute, and when a decision should move to another mode.

2. Fairness and Regulation

A consistent system can also be consistently unfair. The best-known case is the health-care algorithm studied by Obermeyer and colleagues, which used past spending as a proxy for medical need.

Because less had historically been spent on Black patients, the algorithm judged them healthier than they were [14]. No individual case would have exposed the flaw; it took an analysis of outcomes across many patients. That is why the framework builds in sample audits and outcome records.

Regulators have arrived at much the same position. In the United States, the NIST framework organises its guidance under four functions – govern, map, measure and manage [15]. The EU's AI Act imposes heavier obligations on high-risk systems, including human oversight [16]. In both, the level of control rises with the level of risk, which is the logic behind our second question.

3. India

Indian policy has presented AI as a route to inclusive growth since NITI Aayog's 2018 strategy [17]. The India AI Governance Guidelines, released by the Ministry of Electronics and Information Technology in November 2025, set out seven principles: trust, people first, innovation over restraint, fairness and equity, accountability, understandable by design, and safety, resilience and sustainability [18]. They rely on existing law and sector regulators rather than a separate AI Act.

One consequence is easily missed. Principle-based guidance leaves the judgement to the organisation; no regulator will publish a list of decisions that may safely be automated. Our framework gives managers a way to make that call and to explain afterwards how they made it. The requirement of a named owner answers the accountability principle, and the demand for reasons alongside recommendations answers the principle of understandability.

We have small firms, cooperative banks and government offices particularly in mind. Many are adopting AI tools without any governance staff, and the rule they follow has to be simple.

VI. LIMITATIONS, FUTURE RESEARCH AND CONCLUSION

1. Limitations and Future Research

We have already noted that the framework is untested and the review selective. Two further caveats apply. We have treated structure and stakes as high or low, whereas both are matters of degree. And the examples in Table I should not be read too literally: a decision that is routine in a large bank may be anything but routine in a start-up.

Several lines of research follow. The propositions could be tested through surveys and case studies in Indian organisations at different stages of adoption. Longitudinal studies of early-career managers could show whether working with AI strengthens or stunts their judgement. And as AI agents that act with little supervision spread, it will matter to know how quickly decisions drift from

Supervise to Automate once an agent is in place – and whether anyone in the organisation notices.

2. Conclusion

Cheap prediction has not made human judgement redundant. The studies reviewed here show people and machines doing better together on some decisions and worse on others, and they

suggest the pattern is predictable. Routine, low-stakes decisions can be left to software. Routine, high-stakes decisions need a reviewer who brings knowledge of their own and can see the system's reasons. Novel, low-stakes decisions are the place to experiment. Novel, high-stakes decisions belong to accountable people, with AI in support.

This is why we see bridging AI and decision making as a leadership task. It requires leaders to know which decisions their organisation makes, to allocate them deliberately, to keep a person answerable for each, and to review honestly how they turned out. An organisation that does this gains the speed and consistency of the machine while keeping what makes a decision trustworthy: someone who stands behind it.

Acknowledgement

[Optional: acknowledge the institution, guide or funding body. Delete this section if not required.]

REFERENCES

1. H. A. Simon, Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization. New York, NY, USA: Macmillan, 1947.
2. A. Agrawal, J. Gans, and A. Goldfarb, Prediction Machines: The Simple Economics of Artificial Intelligence. Boston, MA, USA: Harvard Business Review Press, 2018.
3. M. Vaccaro, A. Almaatouq, and T. Malone, "When combinations of humans and AI are useful: A systematic review and meta-analysis," Nature Human Behaviour, vol. 8, no. 12, pp. 2293–2303, 2024, doi: 10.1038/s41562-024-02024-1.
4. D. Kahneman, O. Sibony, and C. R. Sunstein, Noise: A Flaw in Human Judgment. New York, NY, USA: Little, Brown Spark, 2021.
5. J. Kleinberg, H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan, "Human decisions and machine predictions," The Quarterly Journal of Economics, vol. 133, no. 1, pp. 237–293, 2018, doi: 10.1093/qje/qjx032.
6. M. H. Jarrahi, "Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making," Business Horizons, vol. 61, no. 4, pp. 577–586, 2018, doi: 10.1016/j.bushor.2018.03.007.
7. Y. R. Shrestha, S. M. Ben-Menahem, and G. von Krogh, "Organizational decision-making structures in the age of artificial intelligence," California Management Review, vol. 61, no. 4, pp. 66–83, 2019, doi: 10.1177/0008125619862257.
8. S. Raisch and S. Krakowski, "Artificial intelligence and management: The automation–augmentation paradox," Academy of Management Review, vol. 46, no. 1, pp. 192–210, 2021, doi: 10.5465/amr.2018.0072.
9. B. J. Dietvorst, J. P. Simmons, and C. Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err," Journal of Experimental Psychology: General, vol. 144, no. 1, pp. 114–126, 2015, doi: 10.1037/xge0000033.
10. J. M. Logg, J. A. Minson, and D. A. Moore, "Algorithm appreciation: People prefer algorithmic to human judgment," Organizational Behavior and Human Decision Processes, vol. 151, pp. 90–103, 2019, doi: 10.1016/j.obhdp.2018.12.005.
11. R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," Human Factors, vol. 39, no. 2, pp. 230–253, 1997, doi: 10.1518/001872097778543886.
12. E. Brynjolfsson, D. Li, and L. R. Raymond, "Generative AI at work," The Quarterly Journal of Economics, vol. 140, no. 2, pp. 889–942, 2025, doi: 10.1093/qje/qjae044.
13. F. Dell'Acqua et al., "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality," Harvard Business School, Boston, MA, USA, Working Paper 24-013, 2023. [Online]. Available: https://ssrn.com/abstract=4573321
14. Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," Science, vol. 366, no. 6464, pp. 447–453, 2019, doi: 10.1126/science.aax2342.
15. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," Gaithersburg, MD, USA, NIST AI 100-1, Jan. 2023, doi: 10.6028/NIST.AI.100-1.
16. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Official Journal of the European Union, L series, 12 Jul. 2024.
17. NITI Aayog, "National Strategy for Artificial Intelligence #AIforAll," Government of India, New Delhi, India, Jun. 2018.
18. Ministry of Electronics and Information Technology, "India AI Governance Guidelines," Government of India, New Delhi, India, Nov. 2025.