International Journal of Economic Practices and Theories (IJEPT)
ISSN: 2247-7225
Volume 2026, Issue 1

Artificial Intelligence-Driven Decision-Making in Modern Business Management

Dr. Manisha V Patil [1], Mr. Anil P. Patil [2], Govind S [3], Dr. Pavan G Bhadang [4], Prashant Pandey [5]

[1] Assistant Professor, Faculty of Education and Humanities; Assistant Director, Career Development and Corporate Relations Cell, Training and Placement Office, JSPM University, Pune, Maharashtra, India.

[2] Sr. Manager, Department of Application Engineering, Schaeffler India Ltd, Pune, Pune District, Maharashtra, India.

[3] Technology Consultant, Emvigo Technologies Pvt Ltd.

[4] Assistant Professor, Department of Management Studies, Yeshwantrao Chavan College of Engineering, Nagpur, India.

[5] Assistant Professor, Department of Commerce, Vignana Jyothi Institute of Arts and Sciences, Hyderabad, India.

Abstract:

Artificial intelligence (AI) is changing how businesses decide: it lowers the cost of prediction, widens the range of decisions that can be supported or automated, and introduces a non-human participant into organizational judgment. Evidence on whether this improves decisions remains scattered across management, information systems and decision science. This paper synthesizes that literature through a structured narrative review and develops a five-layer framework of AI-driven decision-making that links data, models, decision structures, governance and outcomes, together with organizational enablers and a feedback loop. The paper maps AI techniques to operational, tactical and strategic decisions, distinguishes delegated, hybrid and aggregated decision structures, and derives six testable propositions on AI capability, decision-structure fit, the automation–augmentation balance, appropriate reliance, governance and expert participation. Functional illustrations in finance, marketing, operations, human resources and strategy show how the framework applies, and a four-stage road-map translates it into practice. The analysis argues that AI-driven decision-making is an organizational design problem: value depends less on the sophistication of models than on complementary resources, clear decision rights and accountable governance. The paper closes with implications for managers and a research agenda.

Index Terms—artificial intelligence, decision-making, business management, human–AI collaboration, algorithmic governance, decision support, AI capability, digital transformation.

I. Introduction

Managers have always decided under limited information, limited time and limited attention. The classic account of bounded rationality holds that decision makers do not optimize; they search until an alternative looks good enough [1]. Decades of behavioral research have since documented how intuition, heuristics and bias shape managerial choices, sometimes productively and sometimes at real cost [2, 3]. What has changed in the last decade is that the organization now has a second kind of decision participant: software that learns patterns from data, makes predictions at scale, and increasingly recommends or executes actions.

Three developments explain why this matters now. First, the volume and variety of operational data have grown to the point where data-driven management has become a mainstream practice rather than an experiment [4, 5, 6]. Second, advances in machine learning, deep learning and, more recently, large language and foundation models have widened the range of tasks that can be automated or supported, from forecasting and scoring to drafting, summarizing and reasoning over unstructured text [7, 8, 9]. Third, the economics of prediction have shifted: when prediction becomes cheap, the value of the complementary inputs, namely data, judgment and action, rises [10].

Earlier waves of automation reshaped routine physical and clerical work; the present wave reaches cognitive tasks that were long treated as uniquely human [11].

Practitioner enthusiasm has outrun organizational understanding. Surveys of early corporate adopters reported that the most visible barriers were rarely algorithmic; they concerned data, skills, process fit and leadership alignment [12]. Academic work reaches a similar conclusion from a different direction. AI capability is a bundle of tangible, human and intangible resources, not a single tool, and its performance effects depend on how those resources are orchestrated [13]. Multidisciplinary reviews of the field stress that opportunities, challenges and policy questions must be considered together [14]. Meanwhile, scholars have asked who should decide what: when a decision should be delegated to an algorithm, when it should be shared, and when it should remain firmly human [15, 16, 17].

This paper addresses a practical question that sits behind these debates: how should a modern business organize AI-driven decision-making so that decisions become faster and more accurate without becoming opaque, biased or unaccountable? We answer it by integrating three streams of work (decision theory, information systems and capability research, and human–AI interaction) into one framework that managers can use and researchers can test.

The paper makes four contributions. (1) It maps AI techniques to decision types in a way that makes the fit between problem structure and method explicit. (2) It proposes a five-layer framework of AI-driven decision-making that links data, models, decision structures, governance and outcomes, together with the organizational enablers that connect them. (3) It derives six testable propositions that clarify where current evidence is thin. (4) It offers a staged implementation roadmap and a research agenda. The analysis is conceptual and review-based; it does not report new field data, and its propositions are intended to guide empirical work rather than to substitute for it.

The remainder of the paper is organized as follows. Section II reviews the theoretical background. Section III explains the review approach. Section IV classifies AI techniques by decision type. Section V develops the framework and propositions. Section VI discusses functional applications. Section VII presents an illustrative scenario. Section VIII examines risks, trust and governance. Section IX sets out an implementation roadmap. Section X discusses recurring tensions, Section XI draws implications and a research agenda, and Section XII concludes.

II. Theoretical Background

A. Decision-Making in Organizations

Organizational decisions differ in structure, stakes and speed. Strategic decisions are infrequent, ambiguous and consequential, and their quality depends on process as much as on information [18]. Operational decisions, by contrast, are frequent, well specified and measurable, which makes them natural candidates for automation. Between the two lies a broad middle ground of tactical choices such as pricing, staffing, credit limits, inventory positions and campaign allocation, where data are rich but context still matters. A useful way to think about AI is therefore not as a replacement for the decision maker but as a change in the cost and availability of particular decision inputs, above all prediction [10].

Behavioral research gives two reasons to take this seriously. Human judgment is efficient but variable and prone to systematic error [2], and simple, transparent rules can sometimes outperform complex judgment when data are limited [3]. Both observations cut in favor of well-designed algorithmic support. Yet the same research warns that how people respond to algorithms is itself a behavioral phenomenon, a point taken up in Section VIII.

B. From Decision Support Systems to AI-Driven Decisions

Decision support systems (DSS) emerged to help managers structure problems, run models and query data [19]. Business intelligence and analytics extended this idea to enterprise-scale reporting and diagnostics [4]. Modern AI changes the character of the support in three ways. It learns from data rather than following hand-coded rules [7, 20]. It can operate on unstructured inputs such as text, images and speech. And it can act, not merely advise, when connected to workflows and transaction systems. The history of the field also shows cycles of optimism and disappointment, which is a reminder that capability claims should be tested against organizational reality [21, 22].

The management of AI has been described as managing three interdependent facets: autonomy, learning and inscrutability [23]. Autonomy concerns how much the system acts without human intervention; learning concerns how it improves and drifts over time; inscrutability concerns how hard its reasoning is to inspect. Each facet maps to a managerial choice, and together they define the decision problem that Sections V and VIII develop.

C. Capabilities, Resources and Adoption

The resource-based view treats sustained advantage as the result of valuable, rare and hard-to-imitate resources [24], and dynamic-capability theory adds the ability to sense, seize and reconfigure resources as environments shift [25]. Applied to AI, this suggests that data and algorithms alone are weak sources of advantage because both are increasingly available. What differentiates firms is the combination of proprietary data, skilled people, cross-functional coordination and an organizational orientation toward learning from algorithmic output [13, 26]. Big-data work in strategy makes a related point: new data sources change the context of strategic decisions, including what counts as evidence and who holds interpretive authority [27].

Adoption theories complement this view. The technology acceptance model and its unified successor emphasize perceived usefulness, ease of use and social influence [28, 29], while the technology–organization–environment framework adds firm context and external pressure [30]. For AI-driven decisions, both lines of work imply that adoption at the individual level (will the manager use the recommendation?) and at the organizational level (will the firm redesign decision rights?) are distinct problems that must each be solved.

III. Review Approach

The study follows a structured narrative review. The aim was to build a concept-centric synthesis rather than to count publications, following the logic of concept-centered literature reviews in information systems and the principle of transparent, evidence-informed review in management [31, 32]. Sources were drawn from leading peer-reviewed journals in management, information systems and decision science, supplemented by foundational books and by regulatory and standards documents where the topic requires them. Priority was given to works that (i) address AI or algorithmic decision-making in organizations, (ii) offer a theoretical framework or a well-identified empirical result, and (iii) have shaped subsequent debate.

Material was then coded on four dimensions: the type of decision studied, the AI technique or capability involved, the human–AI configuration, and the reported outcome or risk. The coding supports the classification in Section IV and the framework in Section V. Because the review is selective and not exhaustive, its conclusions are framed as propositions. Readers who wish to extend it into a formal systematic review should document database coverage, search strings and screening decisions explicitly.

IV. AI Techniques And Decision Types

AI is not one technology. Decision value depends on matching the technique to the structure of the problem. Table I groups decisions into three levels and links each to the AI approaches that currently fit best, typical uses, and the dominant risk. The grouping is a heuristic, not a law: a single use case, such as credit approval, can involve a predictive model, an optimization routine and a human review step at once.

Table I. AI techniques mapped to decision levels

Decision level	Typical AI approach	Illustrative business uses	Dominant risk
Operational (high volume, well-defined)	Supervised learning, rules engines, optimization, process automation	Demand and fraud scoring, dynamic routing, inventory replenishment, ticket triage	Model drift; silent error at scale

Tactical (moderate volume, partly structured)	Predictive and prescriptive analytics, recommender systems, forecasting ensembles	Pricing, credit limits, workforce scheduling, campaign allocation, vendor selection	Bias in training data; over-reliance
Strategic (low volume, ambiguous)	Scenario simulation, text mining, generative and foundation models as analysts	Market scanning, competitor analysis, option generation, risk narratives	Opacity; false confidence; hallucination

At the operational level, structure and volume favor automation. Decisions are repeated thousands of times, outcomes are observable, and feedback arrives quickly enough to retrain models. The managerial task is mostly about monitoring: detecting drift, setting thresholds and defining escalation paths for exceptions. At the tactical level, predictions improve but do not settle the decision, because the right action depends on costs, constraints and context that the model may not represent. Prescriptive methods can help by combining forecasts with optimization, but a human reviewer commonly remains responsible for the final call.

At the strategic level, AI is better seen as an analyst and a sparring partner than as a decider. Strategic questions involve ambiguity about what the problem is, not only about which option is best. Here, the human–AI relationship has been described as one of complementarity: algorithms handle uncertainty through computation, while people handle complexity and equivocality through interpretation, negotiation and sense-making [16]. Generative models widen this role by drafting scenarios and synthesizing large bodies of text, but their fluency can mask error, so verification routines become part of the decision process rather than an afterthought [9].

A simple, stylized expression captures why the choice among decision structures is economic as well as technical. Let the net value of delegating a decision class to AI be

$$NDV = N \cdot (A_{AI} - A_H) \cdot B - C_{AI} - E[L], \quad (1)$$

where N is decision volume, A_{AI} and A_H are the accuracy of the AI-supported and human-only processes, B is the benefit of a correct decision relative to an incorrect one, C_{AI} is the total cost of building, running and governing the system, and E[L] is the expected loss from rare but severe failures such as bias incidents, regulatory penalties or reputational damage. Equation (1) is not an estimation model; it makes explicit that high volume favors delegation, that accuracy gains must be weighed against governance cost, and that tail risks can reverse an apparently positive case.

V. A Framework For Ai-Driven Decision-Making

A. Framework Overview

Fig. 1 presents the proposed framework. It has five layers that progress from data to outcomes and a set of organizational enablers that act on every layer. A feedback loop returns outcome evidence to the data and model layers, because decision systems that do not learn from their consequences degrade.

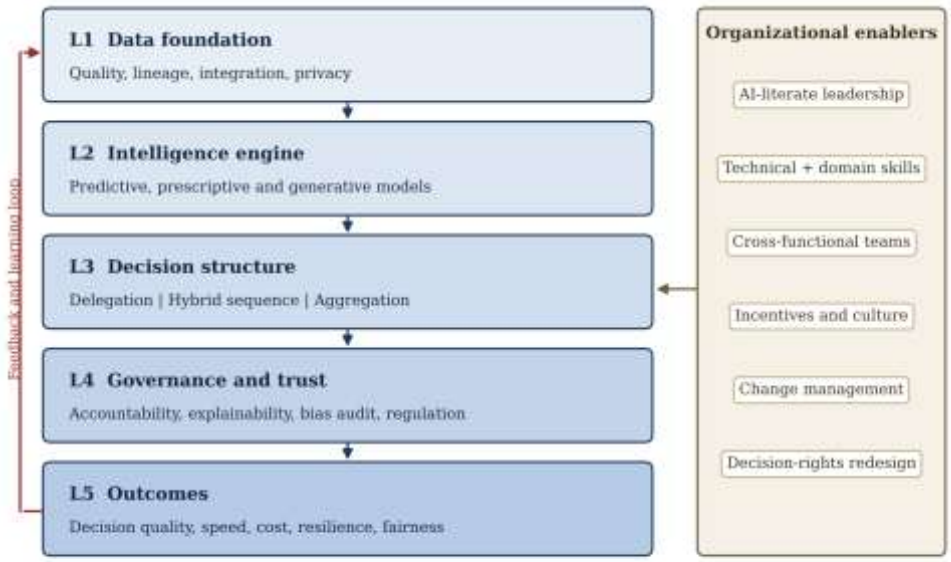


Fig. 1. Five-layer framework of AI-driven decision-making with organizational enablers and a feedback loop.

Layer 1, the data foundation, covers quality, integration, lineage and privacy. Poor data propagate directly into poor decisions, and in practice the largest share of effort in AI programs goes here. Layer 2, the intelligence engine, contains the predictive, prescriptive and generative models that transform data into recommendations. Layer 3, the decision structure, specifies who or what decides. Layer 4, governance and trust, sets the controls (accountability, explainability, bias audit and regulatory compliance) that make the system safe to rely on. Layer 5 measures outcomes along several dimensions, not accuracy alone: decision quality, speed, cost, resilience and fairness.

B. Data and Model Layers in Practice

The first two layers deserve more attention than they usually receive in management writing, because they determine the ceiling of everything above them. On the data side, organizations face four recurring problems. Data are fragmented across systems that were never designed to interoperate. Historical records encode past practice, including its errors and inequities. Key variables are often proxies for what actually matters, as when cost stands in for need or clicks stand in for satisfaction [33]. And the data that would be needed for a new kind of decision frequently do not yet exist. Building a reliable data foundation therefore involves governance, not only engineering: someone must own each critical dataset, define its meaning, and decide who may use it and for what purpose [5, 6].

On the model side, three families of method dominate current business use. Predictive models estimate the probability or magnitude of an uncertain outcome, such as default, churn or demand. Prescriptive models combine predictions with constraints and objectives to recommend an action, such as a production plan or a price schedule. Generative models produce text, code or images and can act as drafting and analysis assistants [8, 9]. Selecting among them is a managerial question about the decision, not only a technical one: a prescriptive model is only as good as the objective function it is given, and a generative model is only as reliable as the verification step that follows it. Model risk therefore belongs in the same register as other operational risks, with owners, limits and review cycles [34].

C. Decision Structures

Layer 3 is where the framework makes its most distinctive claim. Drawing on the contingency view of decision structures, three arrangements can be distinguished: full delegation to the algorithm, hybrid sequential arrangements in which human and algorithm hand work back and forth, and aggregated decisions in which human and machine outputs are combined [15]. The best arrangement depends on how specific the search space is, how interpretable the process must be, how many alternatives exist, how fast the decision must be made, and how replicable the task is. Delegation suits specific, high-speed, replicable problems. Hybrid and aggregated structures suit problems in which interpretability and contextual judgment matter.

The automation–augmentation distinction should not be treated as a clean choice. Evidence and theory suggest that the two are interdependent over time and across the organization: today's augmentation generates the data and routines that enable tomorrow's automation, and heavy automation can erode the skills needed to supervise it [17]. The framework therefore treats the decision structure as a dynamic setting that is reviewed periodically, not fixed at design time.

D. Enablers

The enablers in Fig. 1 reflect what capability research has repeatedly found: organizational and human factors, rather than technical ones, often separate firms that benefit from AI from those that do not [12, 13]. AI-literate leadership sets priorities and absorbs risk. Technical and domain skills must be combined, since models built without domain understanding rarely fit the decision they are meant to support. Cross-functional teams reduce hand-off failure. Incentives and culture determine whether people use, ignore or game algorithmic output. Studies of AI development in practice show that experts and developers negotiate what the system should know and how it should be used, so participation is a design requirement and not a courtesy [35]. Finally, decision rights must be redesigned explicitly; otherwise AI recommendations enter an organization with no agreed rule for who may override them.

E. Propositions

The framework yields six propositions, summarized in Table II. They are stated so that each can be tested with firm-level, decision-level or experimental data.

Table II. Research propositions derived from the framework

No.	Proposition	Primary support
P1	Firms with stronger AI capability, defined as the joint presence of data, technical and managerial resources, realize larger decision-quality and performance gains than firms with comparable AI spending but weaker complementary resources.	[13, 26]
P2	The fit between decision structure (delegation, hybrid, aggregated) and decision characteristics (specificity, interpretability, speed, replicability) positively predicts decision quality.	[15]
P3	Balanced attention to automation and augmentation outperforms exclusive emphasis on either, over time.	[17]
P4	Explanation and calibrated performance feedback increase appropriate reliance on AI recommendations, and reduce both algorithm aversion and over-reliance.	[36, 37, 38]
P5	Formal governance (bias audit, accountability assignment, monitoring) reduces the frequency and severity of adverse algorithmic incidents without eliminating the speed advantage of AI.	[33, 34]
P6	Participation of domain experts in model design increases adoption and decision quality, especially for tactical and strategic decisions.	[16, 35]

VI. Functional Applications

The framework applies across business functions, though the balance among decision levels differs. This section illustrates the pattern without claiming to be exhaustive.

Finance and risk: Credit scoring, fraud detection, anomaly detection in transactions and forecasting are among the most established uses of machine learning in business. They are typically high-volume and measurable, so delegation or tightly supervised hybrid structures dominate. The regulatory burden is also heavy: decisions that affect individuals' access to credit invite requirements on explanation, fairness and record-keeping, which makes Layer 4 central rather than peripheral [39, 40].

Marketing and sales: Recommender systems, segmentation, churn prediction and dynamic pricing translate customer data into tactical decisions. Algorithms are good at estimating response probabilities; managers still decide what the brand should and should not do. A typical hybrid design lets the model allocate budget across segments within boundaries that managers set, with periodic review of whether the optimized behavior still serves long-term customer relationships.

Operations and supply chain: Demand forecasting, inventory policy, routing and predictive maintenance fit the operational and tactical levels. Advantage here comes less from a single model than from integration across planning, procurement and logistics, which again points to cross-functional coordination as an enabling factor [13]. Exceptions, such as supply disruptions that the training data have never seen, require escalation paths to humans who can reason about novel situations.

Human resource management: Screening, scheduling, performance analytics and attrition prediction illustrate both the promise and the hazard of AI-driven decisions. They can reduce administrative load and broaden candidate pools, but they also directly affect people's livelihoods. Research on algorithmic management shows that algorithms at work can function as instruments of monitoring and control, with consequences for autonomy and trust [41]. Field studies of AI built for hiring show that its design involves contested judgments about what good performance means [35]. Employment-related uses are among those treated as high risk under the European Union's AI Act, which imposes obligations such as risk management, data governance, transparency and human oversight [40].

Strategy and corporate planning: Generative models can summarize market reports, draft scenarios and surface weak signals in unstructured data. This is the area where the strongest claims are made and the weakest evidence exists. Following the framework, strategic uses should be treated as aggregated or hybrid: the model widens the option set and tests assumptions, while accountable executives interpret, challenge and decide [16, 22].

VII. Illustrative Scenario: Tactical Pricing Decisions

A concrete, stylized case shows how the layers interact. Consider a mid-sized distributor that sets prices for several thousand products across regional customer groups. Under the existing process, regional managers adjust prices monthly using spreadsheets and experience. The firm considers an AI-supported pricing engine.

Layer 1. The firm first discovers that transaction records, contract terms and competitor information sit in separate systems with inconsistent product codes. A data-owner role is created for the pricing domain and a quality threshold is set before any model is trained. Layer 2. A predictive model estimates demand sensitivity by product and segment, and a prescriptive layer proposes price changes subject to margin floors and contractual constraints. Layer 3. Because the decision is tactical, with moderate volume and meaningful context, the firm chooses a hybrid structure: the engine proposes prices, regional managers may accept or override them within defined bounds, and large changes require a second review. The override rate is monitored.

Layer 4. Governance defines who is accountable for pricing outcomes (the commercial director), requires an audit that checks whether recommendations treat customer groups consistently, and sets rules for explaining a recommendation to a manager in plain terms. Layer 5. The scorecard tracks margin, volume, time spent per pricing cycle, customer complaints and the share of overrides that proved correct. Over several cycles, evidence on overrides tells the firm where the model is weak and where managers add genuine value; low-risk product groups may move toward delegation while strategic accounts stay hybrid. This is the dynamic view of decision structure that the framework recommends, and it shows why P2 and P3 are linked: structure is chosen on fit, then adjusted as evidence accumulates [15, 17].

VIII. Risks, Trust And Governance

A. Human Responses to Algorithms

The benefit of an algorithm depends on whether people use it appropriately. Experiments show that individuals often lose confidence in an algorithm faster than in a human after seeing it make a mistake, even when the algorithm is more accurate overall, a pattern known as algorithm aversion [36]. Other studies find the opposite pattern in many settings: people weight advice more heavily when they believe it comes from an algorithm, which has been called algorithm appreciation [37]. These results are consistent once context is considered. Reliance depends on task type, on feedback about performance and on the user's sense of control. Reviews of trust in AI emphasize that trust is shaped by system characteristics, by the user and by the situation, and that it is dynamic [42].

The managerial objective is therefore appropriate reliance, not maximal trust. Research on automation has long distinguished misuse (over-reliance), disuse (under-reliance) and abuse (deploying automation without regard for human consequences) [43], and has shown that good design aligns trust with the actual capabilities of the system [38]. Calibrated feedback, visible uncertainty and clear override procedures are practical means to that end.

B. Opacity, Bias and Accountability

Opacity has several sources: proprietary secrecy, technical complexity and the mismatch between how machines represent patterns and how people reason [44]. Explainable-AI methods offer partial remedies, but they differ in fidelity and audience, and an explanation that satisfies a data scientist may not help a loan officer or a regulator [45]. Bias is a separate problem. Algorithms trained on historical data can reproduce historical inequity. A well-known study of a health-management algorithm found that using health cost as a proxy for health need systematically disadvantaged Black patients, an example in which the problem lay in the target variable, not in the modeling technique [33]. The wider literature maps the ethical issues as concerns with evidence, fairness, transparency, autonomy, privacy and traceability of responsibility [39].

Accountability is the organizational counterpart of these technical concerns. When an AI-supported decision harms someone, a named human or committee must be answerable. This is why the framework places accountability assignment at Layer 4 and why decision rights are listed as an enabler: unclear ownership is both an ethical hazard and an operational one.

C. Standards and Regulation

Principles for ethical AI have converged around themes such as transparency, justice and fairness, non-maleficence, responsibility and privacy, although implementation differs widely [46, 47]. Two instruments are particularly relevant to business decision-makers. The European Union's AI Act takes a risk-based approach and attaches stricter obligations to high-risk uses [40]. The NIST AI Risk Management Framework offers a voluntary structure organized around governing, mapping, measuring and managing AI risk [34]. Organizations that operate across jurisdictions benefit from adopting a single internal standard that satisfies the most demanding requirement, then documenting local variations.

TABLE III. Governance controls by framework layer

Layer	Control	Evidence of operation
L1 Data	Data quality standards, lineage tracking, privacy-by-design	Data catalog; documented data owners; access logs
L2 Models	Validation before release; drift and performance monitoring	Model cards; periodic back-testing reports
L3 Decisions	Explicit decision rights and override rules	Decision-rights matrix; override log and review
L4 Governance	Bias audit, explainability standard, incident response	Audit reports; incident register; named accountable owner
L5 Outcomes	Multi-dimensional scorecard (quality, speed, cost, fairness)	Quarterly review by an oversight committee

IX. Implementation Roadmap

Implementation succeeds when it is staged. Table IV outlines four stages that move an organization from isolated pilots to embedded, governed decision systems. The stages are cumulative: skipping foundations does not shorten the journey, it moves the cost to later failure.

Table IV. Staged maturity roadmap for AI-driven decision-making

Stage	Focus	Typical actions	Exit criterion
1. Foundations	Data, skills, executive sponsorship	Audit data quality; appoint decision owners; train managers in AI literacy	Reliable data for at least one priority decision
2. Pilot	Narrow, measurable use cases	Select decisions with clear metrics; run human-in-the-loop pilots; benchmark against current practice	Documented improvement over baseline
3. Scale	Process redesign and integration	Embed models in workflows; define override rules; stand up monitoring	Stable performance across units
4. Institutionalize	Governance and learning	Formal oversight; periodic review of decision structure; reskilling	Audit-ready evidence and continuous improvement

Three practical principles cut across the stages. First, begin with the decision, not the model: define who decides, with what information, at what cost of error, and then ask which AI capability would change that decision. Second, measure against the realistic alternative, which is the current human or rule-based process, not against perfection. Third, design for exceptions: every automated path should have a clear route to human review, and every human review should feed back into the data that retrain the system.

Measurement design is a part of implementation that is often neglected. A credible evaluation compares AI-supported decisions with a defined baseline over a comparable period, ideally with randomization across units or time. It separates accuracy gains from process gains such as speed and consistency, and it tracks harms as carefully as benefits, for example differential error rates across customer or employee groups. Reporting should include the cost of governance and monitoring, because ignoring it overstates net value, as Equation (1) indicates. Finally, evaluation should continue after deployment: models drift as customers, competitors and regulations change, and a system that was well calibrated at launch may not remain so.

Change management deserves equal weight with engineering. Employees who fear displacement or surveillance will resist or circumvent systems regardless of their quality [41]. Early involvement of the people who make the decisions today, transparent communication about what the system will and will not do, and investment in reskilling are therefore parts of the technical plan, not additions to it [35].

X. Discussion: Managing Tensions

Taken together, the framework and the review point to four tensions that recur across functions and industries. They are not problems to be solved once but trade-offs to be managed continuously, and Table V summarizes how each might be handled.

Table V. Recurring tensions in AI-driven decision-making

Tension	Why it arises	Management response
Speed vs. scrutiny	Automation delivers decisions in milliseconds; review takes time and attention	Reserve human review for high-stakes or low-confidence cases; sample the rest for audit
Accuracy vs. interpretability	The most accurate models are often the least transparent	Match interpretability requirements to decision stakes; use simpler models where explanation is mandatory
Automation vs. augmentation	Each reinforces the other over time; emphasizing one erodes the other	Review decision structures periodically; protect the skills needed to supervise automation
Standardization vs. discretion	Consistency reduces noise but can ignore context	Define override rules; analyze overrides to learn where discretion adds value

The first tension, speed against scrutiny, follows from the economics in Equation (1): the same property that makes delegation attractive, namely high volume at low marginal cost, also multiplies the effect of any hidden error. The practical response is tiered oversight. Low-stakes, high-confidence decisions can flow through with statistical monitoring, while decisions that are high-stakes, novel or low-confidence are routed to a person. The second tension, accuracy against interpretability, is often overstated in general but real in specific settings. Where the law or the stakes require an explanation to the person affected, a somewhat less accurate but more transparent model may be the better choice, and post-hoc explanation methods should be used with a clear view of their limits [44, 45].

The third and fourth tensions concern people. If automation takes over routine decisions, the pool of practitioners who understand the underlying work may shrink, making it harder to detect model failure; this is the substance of the automation–augmentation paradox [17]. Standardization, meanwhile, brings consistency and reduces the noise that characterizes unaided human judgment, yet it may suppress the local knowledge that managers hold. Analyzing overrides is a way of turning this tension into learning: an override that proves right is evidence about what the model lacks, and one that proves wrong is evidence about where human judgment is least reliable [2, 16].

XI. Implications And Research Agenda

Implications for managers. Treat AI as a decision-design problem. Allocate investment to complementary resources (data quality, skills, process redesign, governance) at least as carefully as to models. Choose decision structures deliberately and revisit them as the system and the organization learn. Define appropriate reliance as a goal, and measure it, for example by tracking the override rate and the accuracy of overrides separately.

Implications for researchers. The propositions in Table II call for several kinds of evidence. Firm-level panels can test P1 by separating AI spending from AI capability. Field experiments inside organizations can test P2 and P3 by randomizing decision structures across comparable units. Laboratory and field studies of explanation and feedback can test P4 while clarifying when aversion and appreciation dominate. Comparative studies of governance regimes can test P5, including the cost side of oversight. Process studies of AI development can deepen P6.

Four open questions deserve priority. How do generative systems change the economics and the risks of strategic decisions, where verification is costly? How should responsibility be allocated when several models, vendors and teams contribute to a single decision? How do reskilling and job redesign evolve when augmentation gradually

becomes automation? And how can small and medium enterprises, which lack the resources of large firms, adopt AI-driven decision-making responsibly? Work that addresses developing-economy contexts, where data infrastructure and regulation differ, would be especially valuable.

Limitations. As a selective conceptual review, the paper cannot claim exhaustive coverage, and the framework and propositions have not been tested on primary data. The pace of change in generative AI means that some applications described here will mature or be superseded quickly. These limitations are the reason the contributions are stated as a framework and testable propositions.

XII. Conclusion

AI-driven decision-making is best understood as an organizational design challenge. Algorithms can lower the cost of prediction and extend the reach of data-driven management, but the value of that capability is realized or lost in the surrounding system of data, skills, decision rights and governance. The five-layer framework offered here connects those elements and shows how decision structure should vary with the specificity, interpretability, speed and replicability of the decision at hand. The six propositions identify where evidence is most needed, and the staged roadmap translates the framework into action.

The overall message for business leaders is neither enthusiasm nor caution but discipline. Define the decision, choose the structure that fits it, build the complementary capabilities, keep a named human accountable, and measure outcomes broadly enough to catch what accuracy alone will miss. Organizations that do so are more likely to obtain the speed and consistency of AI without surrendering the judgment and responsibility on which good management depends.

References

[1] H. A. Simon, "A behavioral model of rational choice," Quarterly Journal of Economics, vol. 69, no. 1, pp. 99–118, 1955.

[2] D. Kahneman, Thinking, Fast and Slow. New York, NY, USA: Farrar, Straus and Giroux, 2011.

[3] G. Gigerenzer and W. Gaissmaier, "Heuristic decision making," Annual Review of Psychology, vol. 62, pp. 451–482, 2011.

[4] A. McAfee and E. Brynjolfsson, "Big data: The management revolution," Harvard Business Review, vol. 90, no. 10, pp. 60–68, 2012.

[5] F. Provost and T. Fawcett, "Data science and its relationship to big data and data-driven decision making," Big Data, vol. 1, no. 1, pp. 51–59, 2013.

[6] E. Brynjolfsson and K. McElheran, "The rapid adoption of data-driven decision-making," American Economic Review, vol. 106, no. 5, pp. 133–139, 2016.

[7] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA, USA: MIT Press, 2016.

[8] T. B. Brown et al., "Language models are few-shot learners," in Proc. Advances in Neural Information Processing Systems (NeurIPS), vol. 33, 2020, pp. 1877–1901.

[9] R. Bommasani et al., "On the opportunities and risks of foundation models," arXiv:2108.07258, 2021.

[10] A. Agrawal, J. Gans, and A. Goldfarb, Prediction Machines: The Simple Economics of Artificial Intelligence. Boston, MA, USA: Harvard Business Review Press, 2018.

[11] E. Brynjolfsson and A. McAfee, The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. New York, NY, USA: W. W. Norton, 2014.

[12] T. H. Davenport and R. Ronanki, "Artificial intelligence for the real world," Harvard Business Review, vol. 96, no. 1, pp. 108–116, 2018.

[13] P. Mikalef and M. Gupta, "Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance," Information & Management, vol. 58, no. 3, Art. no. 103434, 2021, doi: 10.1016/j.im.2021.103434.

[14] Y. K. Dwivedi et al., "Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy," International Journal of Information Management, vol. 57, Art. no. 101994, 2021.

[15] 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.

[16] 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.

[17] 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.

[18] K. M. Eisenhardt and M. J. Zbaracki, "Strategic decision making," Strategic Management Journal, vol. 13, no. S2, pp. 17–37, 1992.

[19] D. J. Power, Decision Support Systems: Concepts and Resources for Managers. Westport, CT, USA: Quorum Books, 2002.

[20] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2021.

[21] M. Haenlein and A. Kaplan, "A brief history of artificial intelligence: On the past, present, and future of artificial intelligence," California Management Review, vol. 61, no. 4, pp. 5–14, 2019.

[22] Y. Duan, J. S. Edwards, and Y. K. Dwivedi, "Artificial intelligence for decision making in the era of Big Data – evolution, challenges and research agenda," International Journal of Information Management, vol. 48, pp. 63–71, 2019, doi: 10.1016/j.ijinfomgt.2019.01.021.

[23] N. Berente, B. Gu, J. Recker, and R. Santhanam, "Managing artificial intelligence," MIS Quarterly, vol. 45, no. 3, pp. 1433–1450, 2021, doi: 10.25300/MISQ/2021/16274.

[24] J. Barney, "Firm resources and sustained competitive advantage," Journal of Management, vol. 17, no. 1, pp. 99–120, 1991.

[25] D. J. Teece, G. Pisano, and A. Shuen, "Dynamic capabilities and strategic management," Strategic Management Journal, vol. 18, no. 7, pp. 509–533, 1997.

[26] S. F. Wamba, A. Gunasekaran, S. Akter, S. J. Ren, R. Dubey, and S. J. Childe, "Big data analytics and firm performance: Effects of dynamic capabilities," Journal of Business Research, vol. 70, pp. 356–365, 2017.

[27] I. D. Constantiou and J. Kallinikos, "New games, new rules: Big data and the changing context of strategy," Journal of Information Technology, vol. 30, no. 1, pp. 44–57, 2015.

[28] F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," MIS Quarterly, vol. 13, no. 3, pp. 319–340, 1989.

[29] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," MIS Quarterly, vol. 27, no. 3, pp. 425–478, 2003.

[30] L. G. Tornatzky and M. Fleischer, The Processes of Technological Innovation. Lexington, MA, USA: Lexington Books, 1990.

[31] J. Webster and R. T. Watson, "Analyzing the past to prepare for the future: Writing a literature review," MIS Quarterly, vol. 26, no. 2, pp. xiii–xxiii, 2002.

[32] D. Tranfield, D. Denyer, and P. Smart, "Towards a methodology for developing evidence-informed management knowledge by means of systematic review," British Journal of Management, vol. 14, no. 3, pp. 207–222, 2003.

[33] 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.

[34] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, Gaithersburg, MD, USA, Jan. 2023.

[35] E. van den Broek, A. Sergeeva, and M. Huysman, "When the machine meets the expert: An ethnography of developing AI for hiring," MIS Quarterly, vol. 45, no. 3, pp. 1557–1580, 2021.

[36] 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.

[37] 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.

[38] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," Human Factors, vol. 46, no. 1, pp. 50–80, 2004.

[39] B. D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, and L. Floridi, "The ethics of algorithms: Mapping the debate," Big Data & Society, vol. 3, no. 2, pp. 1–21, 2016.

[40] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)," Official Journal of the European Union, Jul. 2024.

[41] K. C. Kellogg, M. A. Valentine, and A. Christin, "Algorithms at work: The new contested terrain of control," Academy of Management Annals, vol. 14, no. 1, pp. 366–410, 2020.

[42] E. Glikson and A. W. Woolley, "Human trust in artificial intelligence: Review of empirical research," Academy of Management Annals, vol. 14, no. 2, pp. 627–660, 2020.

[43] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," Human Factors, vol. 39, no. 2, pp. 230–253, 1997.

[44] J. Burrell, "How the machine 'thinks': Understanding opacity in machine learning algorithms," Big Data & Society, vol. 3, no. 1, pp. 1–12, 2016.

[45] A. B. Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020.

[46] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," Nature Machine Intelligence, vol. 1, no. 9, pp. 389–399, 2019.

[47] L. Floridi et al., "AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations," Minds and Machines, vol. 28, no. 4, pp. 689–707, 2018.