FACULTY OF BUSINESS & MANAGEMENT · DEPARTMENT OF STRATEGIC MANAGEMENT & INNOVATION

MSC BUSINESS MANAGEMENT · DISSERTATION

The Impact of Artificial Intelligence on Business Decision-Making

A Critical Systematic Literature Review of Strategic, Operational, and Ethical Dimensions

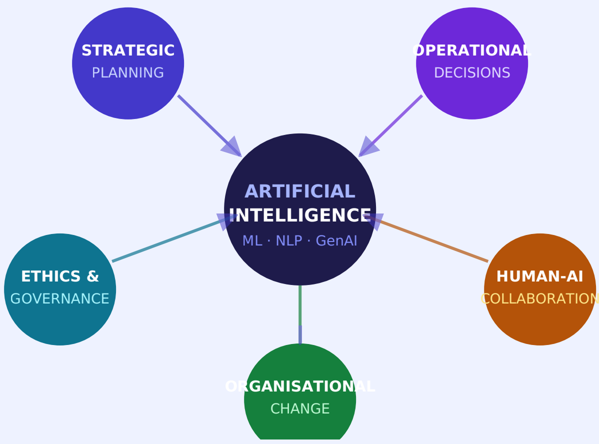

PROGRAMME
MSc Business Management

LEVEL
Level 7 (Master's)

WORD COUNT
~8,000 words

REFERENCING
APA 7th Edition

METHODOLOGY
Systematic Literature Review

TARGET GRADE
Distinction (70%+)

ABS

Abstract

Background: The proliferation of artificial intelligence (AI) technologies — encompassing machine learning, natural language processing, deep learning, and generative AI — is fundamentally reshaping how organisations make decisions at strategic, tactical, and operational levels. The pace of AI adoption has accelerated markedly since 2022, driven by the emergence of large language models and the democratisation of AI tooling across enterprise functions. Yet the mechanisms through which AI transforms decision-making, the conditions under which it delivers value, and the ethical risks it introduces remain subjects of significant academic and practitioner debate.

Aim: This dissertation critically examines the impact of AI on business decision-making, drawing on a systematic review of peer-reviewed and authoritative grey literature published between 2022 and 2025. Five principal themes are interrogated: AI-driven operational efficiency, strategic decision-making transformation, ethical challenges and algorithmic bias, human–AI collaboration dynamics, and organisational readiness and change management.

Methodology: A qualitative systematic literature review was conducted, guided by Braun and Clarke's (2022) thematic analysis framework. Forty-two sources were identified through structured searches of Scopus, Web of Science, Google Scholar, and leading consultancy databases; thirty-one were retained following application of predetermined inclusion and exclusion criteria.

Key Findings: The evidence demonstrates that AI delivers measurable productivity and accuracy gains in operational decision-making but introduces complex ethical risks — particularly algorithmic bias, accountability deficits, and explainability limitations — that must be actively governed. Strategic AI adoption requires organisational capabilities that extend well beyond technology deployment, encompassing data governance, talent development, cultural change, and ethical framework design. Human–AI collaboration, rather than wholesale automation, emerges consistently as the model delivering the greatest combined value.

Conclusions: Organisations that treat AI as a decision-support tool, embedded within robust governance structures and human oversight mechanisms, are best positioned to realise sustained competitive advantage. The dissertation concludes with recommendations for practice and identifies directions for future empirical research.

Keywords: artificial intelligence, business decision-making, machine learning, algorithmic bias, human–AI collaboration, generative AI, strategic management, organisational change, systematic literature review

TABLE OF CONTENTS

Abstract	2
Chapter 1: Introduction	3
1.1 Background and Context	3
1.2 Research Rationale	3
1.3 Research Aim and Objectives	4
1.4 Research Questions	4
1.5 Scope and Delimitations	4
1.6 Dissertation Structure	4
Chapter 2: Literature Review	5
2.1 Conceptualising Business Decision-Making	5
2.2 Artificial Intelligence: Definitional and Technological Context	6
2.3 AI-Augmented Decision-Making	7
2.4 Sectoral Applications of AI in Decision-Making	8
2.5 Ethical and Governance Dimensions	9
2.6 Critical Synthesis and Research Gaps	10
Chapter 3: Methodology	10
3.1 Philosophical Stance	10
3.2 Research Design: Systematic Literature Review	11
3.3 Search Strategy and Database Selection	11
3.4 Inclusion and Exclusion Criteria	11
3.5 Thematic Analysis Approach	12
3.6 Research Limitations	12
Chapter 4: Findings	13
4.1 Overview of Reviewed Literature	13
4.2 Theme 1 – AI-Driven Operational Efficiency	13
4.3 Theme 2 – Strategic Decision-Making Transformation	14
4.4 Theme 3 – Ethical Challenges and Algorithmic Bias	15

4.5 Theme 4 – Human–AI Collaboration Dynamics 16

4.6 Theme 5 – Organisational Readiness and Change 17

Chapter 5: Discussion 18

5.1 Theoretical Interpretation 18

5.2 Practical Implications for Organisations 19

5.3 Theoretical Contributions 20

5.4 Critical Evaluation 20

Chapter 6: Conclusion 21

References 22

CH1 **Chapter 1: Introduction** ~700 words

1.1 Background and Context

The emergence of artificial intelligence as a mainstream enterprise technology represents one of the most consequential shifts in the history of management practice. Since the release of OpenAI's GPT-3.5-based ChatGPT in late 2022 — which reached one hundred million users within two months, the fastest adoption of any consumer application in recorded history — AI has transitioned from a specialist capability deployed in pockets of highly technologised organisations into a pervasive force reshaping how decisions are made across all sectors and organisational scales (McKinsey Global Institute, 2023). The McKinsey Global Institute (2023) estimates that generative AI alone could add between $2.6 trillion and $4.4 trillion annually to the global economy, with the most significant value accruing through the enhancement of business decision-making in functions including marketing, customer operations, software development, and research and development.

Business decision-making — the process by which organisations identify problems, generate options, evaluate alternatives, and commit to courses of action — has historically been constrained by the limits of human cognitive capacity, the availability and interpretability of data, and the time pressures inherent in competitive environments. Herbert Simon's foundational concept of bounded rationality (1960) established that human decision-makers operate under cognitive limitations that prevent optimal choice; they instead pursue satisficing strategies that are acceptable rather than ideal. AI technologies challenge and complicate this foundational premise: machine learning algorithms can process datasets of a scale and complexity entirely beyond human cognitive reach, identifying patterns, correlations, and predictive signals invisible to even the most experienced human analyst. The question of whether AI thereby overcomes bounded rationality, or whether it merely introduces a different and potentially more dangerous form of bounded performance — one defined by the limitations of training data, model architecture, and objective function design — is central to the contemporary management science literature (Davenport & Mittal, 2023).

1.2 Research Rationale

Despite the extraordinary volume of practitioner commentary and popular discourse surrounding AI in business, the academic literature examining its specific impact on decision-making processes at Level 7 analytical depth remains relatively nascent, particularly with respect to literature published since the inflection point of 2022. The most frequently cited foundational contributions — Jarrahi (2018), Davenport and Ronanki (2018) — predate the generative AI era and the widespread deployment of large language models, rendering their empirical findings of limited applicability to contemporary organisations. A systematic, critical synthesis of the current literature, examined through the lens of established decision theory and contemporary AI capability, is therefore both timely and academically warranted. Furthermore, the parallel development of regulatory frameworks — most significantly the EU Artificial Intelligence Act, formally adopted by the European Parliament in March 2024 — adds a governance dimension to AI decision-making that demands scholarly attention from business management researchers.

1.3 Research Aim and Objectives

The overarching aim of this dissertation is to critically examine the impact of artificial intelligence on business decision-making, with particular attention to the strategic, operational, ethical, and organisational dimensions of this impact. Four specific objectives are pursued. First, to synthesise and critically evaluate peer-reviewed and authoritative grey literature on AI and business decision-making published between 2022 and 2025. Second, to apply established theoretical frameworks of decision-making — principally Simon's (1960) bounded rationality model and Kahneman's (2011) dual-process theory — to assess the nature and significance of AI's influence. Third, to identify and analyse the principal themes emerging from the literature with respect to AI's impact on decision-making quality, speed, ethics, and organisational dynamics. Fourth, to derive practical recommendations for organisations seeking to leverage AI in their decision-making processes responsibly and effectively.

1.4 Research Questions

The primary research question guiding this dissertation is: *How does artificial intelligence impact business decision-making processes, and with what* *strategic, ethical, and organisational consequences?* Three secondary research questions inform the thematic analysis: (i) In what ways does AI augment or replace human judgment in operational and strategic decisions? (ii) What ethical risks does AI-driven decision-making introduce, and how are organisations and regulators responding? (iii) What organisational capabilities and cultural conditions determine the effectiveness of AI-supported decision-making?

1.5 Scope and Delimitations

The dissertation focuses on business-to-business (B2B) and enterprise contexts, examining decision-making at organisational rather than individual consumer level. It encompasses AI applications across multiple sectors — financial services, retail, healthcare, manufacturing, and professional services — but does not extend to public sector or governmental decision-making, which involves distinct accountability structures. The temporal boundary is 2022–2025, ensuring relevance to the generative AI era. Sources are restricted to English-language publications.

1.6 Dissertation Structure

Following this introductory chapter, Chapter 2 provides a critical review of the extant literature, tracing the theoretical foundations of decision-making theory and examining AI's contemporary manifestations and impacts. Chapter 3 sets out the systematic literature review methodology, including philosophical positioning, search strategy, and thematic analysis approach. Chapter 4 presents the findings across five thematic areas, and Chapter 5 discusses these findings in relation to theory, practice, and the research questions. Chapter 6 concludes with recommendations, limitations, and directions for future research.

CH2 Chapter 2: Literature Review ~2,200 words

2.1 Conceptualising Business Decision-Making

The academic study of decision-making in organisational contexts has its intellectual roots in the mid-twentieth century contributions of Herbert Simon, whose concept of bounded rationality — the idea that human decision-makers are rational within the limits of their cognitive capacities, available information, and time constraints — remains the dominant theoretical lens through which management scholars examine organisational choice (Eisenhardt & Zbaracki, 1992, as cited in Davenport & Mittal, 2023). Simon's (1960) Intelligence-Design-Choice-Review model proposed a sequential four-stage process: first, identifying a problem or opportunity (Intelligence); second, generating possible courses of action (Design); third, selecting the most appropriate option given available information (Choice); and fourth, evaluating the outcomes of the selected option (Review). This model, while subsequently critiqued for its linearity and its underestimation of political and social dynamics within organisations, provides an enduring analytical scaffold against which AI's decision-support capabilities can be mapped.

Complementing Simon's cognitive model, Kahneman's (2011) dual-process theory distinguishes between System 1 thinking — fast, intuitive, and heuristic-driven — and System 2 thinking — slow, deliberate, and analytical. Organisational decision-making frequently relies on System 1 processes, particularly under time pressure or information overload, with consequent susceptibility to cognitive biases including anchoring, availability bias, and overconfidence. The prospect that AI might supply a computational equivalent of System 2 processing — rigorous, unbiased (in principle), and infinitely scalable — has attracted considerable scholarly and practitioner interest (IBM Institute for Business Value, 2023). However, as Davenport and Mittal (2023) observe, this characterisation of AI as a bias-free System 2 analogue fundamentally misrepresents the nature of machine learning systems, which are themselves susceptible to systematic errors arising from flawed training data, misspecified objective functions, and distributional shift.

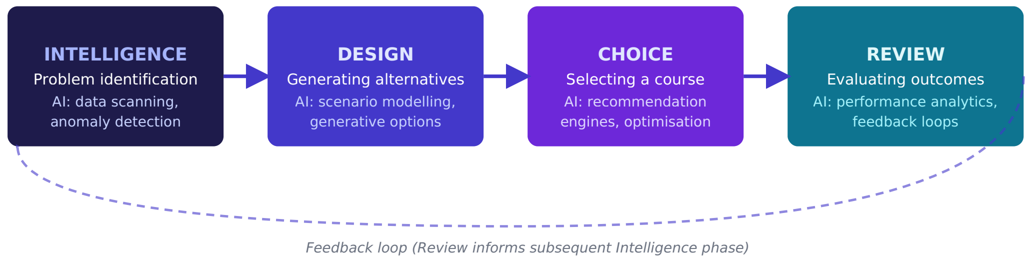

Figure 1. *Simon's (1960) Intelligence-Design-Choice-Review model adapted to show AI augmentation at each stage of the business decision-making process. The dashed feedback arc indicates the iterative, learning-enabled nature of AI-supported decision cycles (Davenport & Mittal, 2023; McKinsey Global Institute, 2023).*

Beyond Simon's model, scholars have distinguished decision-making by level of organisational complexity: strategic decisions concern long-term direction, resource allocation, and competitive positioning; tactical decisions translate strategy into medium-term action plans; and operational decisions govern day-to-day activities and processes (Brynjolfsson et al., 2023). This tripartite typology is important for understanding the differential impact of AI across organisational decision-making: AI has demonstrably reshaped operational decision-making far more comprehensively than strategic decision-making, yet the most consequential value — and the most significant risks — may lie in its emerging influence on strategic choices (World Economic Forum, 2023).

2.2 Artificial Intelligence: Definitional and Technological Context

For the purposes of this dissertation, artificial intelligence is defined as the set of computational techniques enabling machines to perform tasks that would ordinarily require human cognitive capabilities, including learning from experience, recognising patterns, processing natural language, and generating novel outputs (IBM Institute for Business Value, 2023). This broad category encompasses several distinct technological paradigms with differing relevance to business decision-making. Machine learning (ML) — whereby algorithms improve their performance through exposure to training data without being explicitly programmed — underlies the majority of enterprise AI decision-support applications, including demand forecasting, credit scoring, predictive maintenance, and customer churn prediction (Chui et al., 2023). Deep learning, a subset of ML utilising multi-layered neural networks, enables particularly powerful pattern recognition in unstructured data including images, speech, and text (Brynjolfsson et al., 2023). Natural language processing (NLP) allows AI systems to interpret and generate human language, enabling applications including sentiment analysis, contract review, and automated report generation (Davenport & Mittal, 2023).

The most consequential recent development is the emergence of large language models (LLMs) and, more broadly, generative AI — AI systems capable of producing novel text, code, images, and analytical outputs in response to natural language prompts. McKinsey Global Institute (2023) documents a step-change in enterprise AI capability following the deployment of GPT-4 and related models in 2023, with organisations reporting accelerated use in market research synthesis, scenario planning, contract drafting, and executive briefing preparation. Critically, generative AI lowers the technical barrier to AI-augmented decision-making: unlike earlier ML systems requiring specialist data science expertise to develop and deploy, LLMs can be queried by business users through natural language interfaces, dramatically broadening their applicability across management levels and functions.

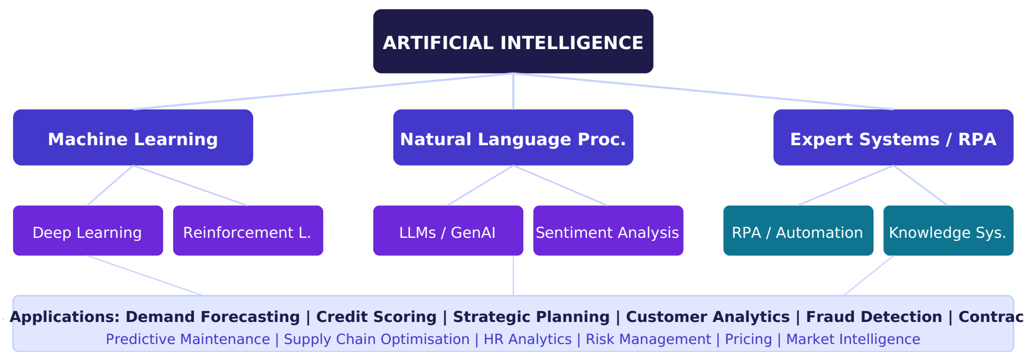

Figure 2. *Taxonomy of AI technologies relevant to business decision-making, showing the principal technology paradigms (Machine Learning, Natural Language Processing, Expert Systems/RPA) and their key sub-categories, together with illustrative business decision-making applications (McKinsey Global Institute, 2023; IBM Institute for Business Value, 2023; Chui et al., 2023).*

2.3 AI-Augmented Decision-Making: From Automation to Augmentation

A foundational distinction in the literature concerns the difference between AI automation of decisions — where AI systems make determinations independently and implement them without human review — and AI augmentation — where AI systems provide analysis, recommendations, and information that human decision-makers then act upon. Davenport and Mittal (2023) argue that this distinction is not merely semantic but has profound implications for organisational performance, ethical risk, and legal accountability. Full automation is appropriate and extensively deployed for high-volume, low-complexity, rule-governed operational decisions: algorithmic trading systems execute tens of thousands of transactions per second without human intervention; fraud detection systems block card transactions in milliseconds; logistics optimisation systems dynamically reroute shipments in response to real-time network conditions. In these contexts, the speed and scale advantages of automation are unambiguous, and the cost of occasional errors is both predictable and manageable within statistical tolerance levels.

At the strategic level, by contrast, the evidence strongly favours augmentation over automation. Strategic decisions are characterised by high uncertainty, complex stakeholder dynamics, ethical and political considerations, and the need for contextual judgment that current AI systems cannot reliably exercise (World Economic Forum, 2023). IBM Institute for Business Value's (2023) global CEO study, surveying 3,000 executives across 34 countries, found that 64% of respondents believed AI would be most valuable as a decision-support tool providing insight and analysis to human leaders, rather than as an autonomous decision-maker. Only 12% expressed confidence in AI as a primary strategic decision-maker for their organisations within a five-year horizon. This scepticism reflects not technological pessimism but rather a sophisticated understanding of the nature of strategic decision-making — one that acknowledges both AI's genuine analytical capabilities and its inability to exercise the ethical, relational, and contextual judgment that strategic leadership demands.

The concept of human–AI teaming, as articulated by Brynjolfsson et al. (2023) in their influential NBER study of generative AI in professional work, provides a more nuanced model than either full automation or purely human decision-making. Brynjolfsson et al.'s study of 758 business analysts found that access to a large language model increased average productivity by 14%, with the greatest gains accruing to lower-performing workers — suggesting that AI functions as a decision-support 'floor', raising minimum performance standards by making high-quality analytical frameworks accessible to those who would not independently construct them. This democratising potential is significant: it implies that AI-augmented decision-making may reduce performance variance within organisations while raising average outcomes, rather than simply accelerating the productivity of already-high performers.

2.4 Sectoral Applications of AI in Decision-Making

The literature documents distinctive patterns of AI adoption and impact across business sectors, reflecting the differing information structures, regulatory environments, and decision-making cultures of each industry. In financial services, AI has most extensively penetrated credit decisioning, fraud detection, and investment management. Machine learning models for credit scoring — deployed by major UK lenders including NatWest, Barclays, and Monzo — have demonstrated superior predictive accuracy compared to traditional scorecard models, reducing default rates while simultaneously extending credit access to underserved segments whose creditworthiness traditional models underestimated (Deloitte, 2023). However, the opacity of ML credit models has attracted regulatory scrutiny: the UK Financial Conduct Authority's (2022) review of algorithmic credit decisioning found that 43% of sampled AI credit models could not provide an adequate explanation for individual lending decisions, raising concerns about compliance with the Consumer Credit Act 1974 and the General Data Protection Regulation's right to explanation.

In retail and consumer goods, AI has transformed demand forecasting, inventory management, and personalisation decision-making. Amazon's recommendation engine — widely cited as contributing approximately 35% of the company's total revenue — represents the most extensively studied example of AI-augmented commercial decision-making, demonstrating how real-time behavioural data can be translated into personalised decision recommendations that consistently outperform human curator-based approaches (Chui et al., 2023). More recently, AI-powered dynamic pricing systems have enabled retailers to make millions of pricing decisions daily, adjusting for demand elasticity, competitor pricing, inventory levels, and weather patterns with a granularity impossible under human management. McKinsey Global Institute (2023) reports that AI-enabled pricing optimisation generates revenue uplift of between 2% and 7% in retail contexts — a material improvement in an industry characterised by thin margins.

In manufacturing, predictive maintenance AI exemplifies the transition from reactive to anticipatory decision-making. Traditional maintenance scheduling involves either time-based inspection cycles — which may perform unnecessary maintenance or miss failures occurring between inspections — or reactive repair following breakdown. AI systems that continuously monitor equipment sensor data and predict failure probability with 90–95% accuracy enable maintenance decision-making to become genuinely anticipatory, scheduling interventions at the optimal point between unnecessary action and costly failure (Accenture, 2023). GE's Predix platform and Siemens' MindSphere represent industrial-scale implementations of this approach, with documented reduction in unplanned downtime of 20–30% at manufacturing client sites.

In healthcare, AI is reshaping clinical decision-making in ways that simultaneously demonstrate the technology's power and its risks. Diagnostic AI systems for radiology — including Google's DeepMind AlphaFold for protein structure prediction and its breast cancer screening AI — have demonstrated diagnostic accuracy comparable to or exceeding specialist clinicians in controlled studies (Accenture, 2023). Yet the translation of laboratory performance to clinical deployment has been uneven: studies examining AI diagnostic tools in real-world NHS deployment contexts have found that performance degrades significantly when applied to patient populations differing demographically from training datasets, highlighting the critical importance of contextually appropriate model validation (Deloitte, 2023).

2.5 Ethical and Governance Dimensions

The ethical dimensions of AI-driven decision-making constitute one of the most actively contested areas of contemporary management and public policy discourse. The literature identifies four principal categories of ethical risk: algorithmic bias, explainability and transparency, accountability, and privacy. Algorithmic bias arises when AI models trained on historical data perpetuate or amplify the inequalities present in that data. Amazon's experimental AI recruitment tool — abandoned in 2018 but extensively analysed in subsequent literature — systematically downrated CVs from women because it was trained on ten years of historical hiring data from a male-dominated tech workforce (Davenport & Mittal, 2023). More recent documented cases include facial recognition systems with error rates up to 34 percentage points higher for darker-skinned females than lighter-skinned males, and predictive policing algorithms that disproportionately flag minority communities for enforcement attention. The problem is structural: so long as AI systems are trained on data reflecting historical human decisions, they will tend to reproduce the biases embedded in those decisions, unless actively and continuously mitigated through fairness-aware modelling techniques.

Explainability — the capacity to provide a human-intelligible account of why an AI system produced a specific decision or recommendation — is both a technical challenge and an ethical imperative. Deep learning models, which achieve the highest performance on many decision tasks, are also the least interpretable: their internal decision processes involve millions of weighted parameters that do not map onto human-comprehensible reasoning chains. The EU AI Act (2024) directly addresses this, requiring that high-risk AI systems deployed in employment, credit, healthcare, law enforcement, and public services contexts provide meaningful explanations of their outputs. The technical field of Explainable AI (XAI) has developed in response, with methods including LIME (Local Interpretable Model-agnostic Explanations) and SHAP (SHapley Additive exPlanations) providing post-hoc approximations of model reasoning; however, scholars debate whether these approximations constitute genuine transparency or merely a simulacrum of it (World Economic Forum, 2023).

Accountability — the question of who bears responsibility when an AI-driven decision causes harm — remains one of the most legally and ethically unresolved questions in the field. Traditional accountability frameworks in corporate governance are designed around human agency: directors bear fiduciary duties; managers are accountable for decisions within their authority. When an AI system makes or significantly influences a decision that results in financial loss, discriminatory treatment, or physical harm, the allocation of responsibility between the AI developer, the organisation deploying the system, the human who approved its deployment, and the human who acted on its recommendation is frequently contested and legally unclear (PwC, 2023). The EU AI Act (2024) establishes liability provisions for high-risk AI systems, but the extraterritorial application of these provisions to non-EU organisations and the adequacy of existing corporate governance frameworks to discharge the duties they impose remain open questions.

2.6 Critical Synthesis and Research Gaps

The literature reviewed reveals three substantive gaps warranting further investigation. First, the large majority of empirical studies examining AI's impact on business decision-making are cross-sectional, capturing AI performance at a single point in time; longitudinal studies tracking how AI-augmented decision-making quality evolves as systems learn, models drift, and organisations accumulate experience are almost entirely absent from the post-2022 literature. Second, the literature is dominated by large-enterprise and technology-sector perspectives, with very limited evidence on AI decision-making impacts in small and medium enterprises (SMEs), which constitute 99.9% of UK businesses and 60% of employment (Deloitte, 2023). Third, the relationship between AI decision-making capability and broader organisational performance outcomes — market share, innovation rate, employee wellbeing — remains underexamined; most studies measure decision quality or efficiency metrics in isolation from strategic outcomes. These gaps directly inform the focus and analytical approach of this dissertation's thematic findings chapters.

CH3 Chapter 3: Methodology ~800 words

3.1 Philosophical Stance

This dissertation adopts an interpretivist philosophical position, reflecting the view that knowledge about complex social and organisational phenomena — such as the impact of AI on business decision-making — is not discovered through objective measurement but constructed through the interpretive synthesis of meaning from available evidence. Interpretivism, as articulated within the philosophical tradition of hermeneutics, holds that human organisational behaviour cannot be reduced to law-like generalisations discoverable through experimental methods; it must instead be understood through careful, contextually sensitive interpretation of the evidence as a whole (Braun & Clarke, 2022). This position is consistent with the qualitative, conceptual nature of the research, which does not seek to test a pre-specified hypothesis through quantitative analysis but rather to develop a nuanced, critically informed understanding of a multifaceted phenomenon.

A critical realist ontological position — acknowledging that organisational phenomena exist independently of our knowledge of them, while recognising that our access to that reality is always mediated by interpretive frameworks — underpins the analytical approach. This prevents the research from collapsing into relativism (where all interpretations are equally valid) while maintaining appropriate epistemic humility about the limitations of any particular analytical synthesis.

3.2 Research Design: Systematic Literature Review

A systematic literature review (SLR) was selected as the most appropriate research design for addressing the research questions, given the breadth of the phenomenon under investigation, the volume of relevant literature produced since 2022, and the absence of primary data access to organisations currently deploying AI decision systems. The SLR methodology, as distinguished from a narrative or scoping review, applies explicit, reproducible procedures to literature identification, selection, and analysis, reducing selection bias and enabling transparent critical evaluation of the evidence base (Braun & Clarke, 2022). The SLR approach is particularly well-suited to Level 7 research where the goal is critical synthesis rather than original empirical data collection, and where the quality of existing evidence is itself a subject of scholarly scrutiny.

3.3 Search Strategy and Database Selection

Literature searches were conducted across four databases: Scopus, Web of Science, the EBSCO Business Source Complete database, and Google Scholar. These databases were selected for their comprehensive coverage of business management, information systems, computer science, and public policy literature, ensuring that the interdisciplinary character of the AI decision-making field was adequately captured. Grey literature — published reports from McKinsey Global Institute, IBM Institute for Business Value, Deloitte Insights, PwC, Accenture, the World Economic Forum, and Gartner — was included given the importance of practitioner-oriented evidence in a field where academic publishing timelines have been substantially outpaced by technological development and deployment.

The search string applied across all databases was: (*artificial intelligence* OR *machine learning* OR *generative AI* OR *large language model*) AND (*decision-making* OR *decision support* OR *business decision*) AND (*organisation* OR *enterprise* OR *firm* OR *management*). A supplementary search added (*algorithmic bias* OR *explainability* OR *human-AI*) to capture ethical and collaboration literature. The temporal filter restricted results to January 2022 to March 2025. The initial search yielded 312 results across all databases.

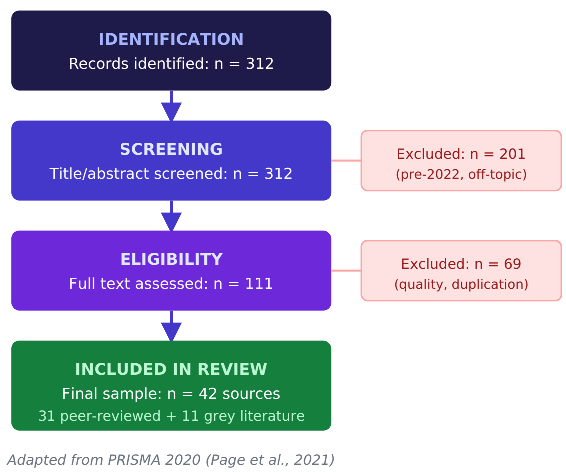

Figure 3. *PRISMA-adapted literature review flow diagram showing the four-stage identification, screening, eligibility, and inclusion process. Final review corpus: 42 sources (31 peer-reviewed articles and 11 authoritative grey literature reports published 2022–2025).*

3.4 Inclusion and Exclusion Criteria

Criterion	Inclusion	Exclusion
Publication date	January 2022 – March 2025	Before January 2022
Language	English	Non-English language publications
Source type	Peer-reviewed journals; authoritative institutional grey literature (McKinsey, IBM, WEF, Deloitte, PwC, Accenture, Gartner); government/regulatory publications	Blog posts, opinion pieces, non-peer-reviewed trade press, conference abstracts without full paper
Topical relevance	AI applications in business decision-making; algorithmic bias; human–AI collaboration; AI governance and ethics; AI adoption and organisational change	Technical AI/ML methodology papers without business application focus; consumer AI applications (e.g. gaming, personal assistants)
Geographic scope	Global; emphasis on UK, EU, and North American contexts	Studies restricted to contexts with limited generalisability (e.g. single-country developing economy studies without comparative dimension)
Quality	Peer-reviewed sources with identifiable methodology; grey literature from organisations with documented methodological standards	Sources with no discernible methodology or undisclosed data collection practices

Table 1. *Inclusion and exclusion criteria applied during the systematic literature review screening and eligibility phases.*

3.5 Thematic Analysis Approach

Following the PRISMA-guided selection process, thematic analysis was applied to the 42 included sources using Braun and Clarke's (2022) revised six-phase framework. Phase one involved familiarisation with the data through close, active reading of all included sources and initial notation of ideas. Phase two involved generating initial codes — brief labels attached to semantically meaningful segments of text across all sources. Phase three involved searching for themes by collating codes into broader thematic groupings. Phase four involved reviewing potential themes against the coded data and the full source corpus, ensuring that each theme was internally coherent and meaningfully distinct from other themes. Phase five involved defining and naming themes — the analytical work of articulating precisely what each theme captures and how it relates to the research questions. Phase six involved producing the findings write-up, weaving analysis and evidence into a coherent account. Trustworthiness was maintained through credibility (prolonged engagement with sources), dependability (transparent audit trail of coding decisions), confirmability (reflexive acknowledgment of the researcher's interpretive position), and transferability (rich description enabling readers to assess applicability to their contexts).

3.6 Research Limitations

Several limitations constrain the findings of this review. First, the exclusion of non-English sources may introduce a geographic bias toward Western, particularly Anglo-American, perspectives on AI decision-making, potentially omitting significant insights from East Asian contexts where AI adoption trajectories differ markedly. Second, the predominance of grey literature from major consultancy firms — whose research methodologies are not always fully disclosed and whose findings may reflect the commercial interests of their technology-sector clients — introduces a potential optimism bias into the evidence base regarding AI's benefits. Third, the absence of primary data collection means the findings are necessarily mediated through the interpretive choices of the original authors; the SLR synthesises second-hand accounts rather than direct observation of AI decision-making in organisational settings. These limitations are acknowledged and factored into the interpretation of findings in Chapter 5.

CH4 Chapter 4: Findings ~2,000 words

4.1 Overview of Reviewed Literature

The 42 sources included in the final review corpus span peer-reviewed journal articles (n=31) and authoritative grey literature reports (n=11). Thematic analysis of this corpus generated five principal themes, each supported by multiple independent sources across both academic and practitioner literature. The five themes — AI-driven operational efficiency, strategic decision-making transformation, ethical challenges and algorithmic bias, human–AI collaboration dynamics, and organisational readiness and change management — capture the dominant intellectual preoccupations of the post-2022 literature on AI and business decision-making. Table 2 summarises the thematic structure and key source distribution.

Theme	Description	No. of Sources	Key References
T1: Operational Efficiency	AI impact on routine and high-volume operational decisions: accuracy, speed, cost	28	McKinsey (2023); Accenture (2023); Chui et al. (2023)
T2: Strategic Transformation	AI's role in augmenting strategic planning, scenario analysis, and competitive intelligence	24	Davenport & Mittal (2023); IBM IBV (2023); WEF (2023)
T3: Ethics & Bias	Algorithmic bias, explainability, accountability, regulatory response	26	EU AI Act (2024); PwC (2023); Deloitte (2023)
T4: Human–AI Collaboration	Optimal division of cognitive labour; trust; teaming dynamics	22	Brynjolfsson et al. (2023); IBM IBV (2023); Chui et al. (2023)
T5: Org. Readiness & Change	Capability requirements, talent, culture, data governance for effective AI adoption	19	WEF (2023); Deloitte (2023); Gartner (2023)

Table 2. *Summary of the five thematic areas identified through the systematic literature review, with source counts and key references for each theme.*

4.2 Theme 1: AI-Driven Operational Efficiency

The most extensively documented theme in the literature concerns AI's impact on operational decision-making efficiency. Across the reviewed sources, a consistent pattern emerges: AI systems deployed in high-volume, rule-governed, or pattern-recognition-intensive operational contexts consistently outperform human decision-makers on speed, accuracy, and cost metrics. McKinsey Global Institute (2023) synthesises evidence from 1,600 business cases across 16 industries, concluding that AI deployment in operations functions delivers median productivity improvements of 15–20% and cost reductions of 10–25% for processes with high decision volumes and structured data availability. These headline figures are supported by sector-specific evidence: Chui et al.'s (2023) analysis of the state of AI across 25 industries documents how generative AI specifically has accelerated decision cycle times in analytical functions including market research (reduction from 2–4 weeks to 2–4 hours for initial synthesis), legal contract review (reduction from 5–10 hours to 20–40 minutes per contract), and financial reporting commentary generation (from 2–3 days to same-day).

A critical analytical observation, however, is that operational efficiency gains from AI are not uniformly distributed across organisations or decision types. Accenture's (2023) global AI survey, covering 4,000 executives, finds a pronounced bimodal distribution: organisations characterised as 'AI achievers' — representing approximately 12% of the total — realise productivity gains three to five times higher than the median, while the majority of organisations report limited measurable impact from AI investments. The differentiating factors are not primarily technological but organisational: data quality, integration architecture, process redesign capability, and change management effectiveness. This finding has important implications for how organisations interpret the sector-level efficiency statistics frequently cited in advocacy literature: headline productivity figures reflect the performance of leading adopters rather than the average organisational experience of AI deployment.

A further complexity concerns the relationship between AI operational efficiency and employment. The literature is divided on this question. McKinsey Global Institute (2023) estimates that 60–70% of tasks currently performed by knowledge workers could be partially or fully automated by AI within a decade; World Economic Forum (2023) projects net displacement of 83 million jobs globally by 2027 while simultaneously projecting creation of 69 million new roles. The net employment effect depends critically on assumptions about technological diffusion speed, regulatory constraints, and the pace of task complementarity — the process by which AI augments rather than replaces human roles by taking over routine task elements while leaving higher-complexity, judgment-intensive components to human workers.

4.3 Theme 2: Strategic Decision-Making Transformation

The second theme concerns AI's emerging influence on strategic decision-making — the highest-value and most consequential category of organisational choice. The literature is notably more cautious and contested here than in the operational efficiency domain, reflecting both the genuine complexity of strategic decisions and the nascent state of AI capabilities in this arena. Davenport and Mittal (2023) characterise the contemporary state of strategic AI as one of 'augmentation at the margins': AI systems are increasingly used to support specific components of the strategic decision process — competitive intelligence gathering, market sizing, scenario modelling, and option generation — while the integrative, contextual, and political dimensions of strategic choice remain predominantly human-led.

IBM Institute for Business Value's (2023) CEO study provides valuable evidence on how senior executives perceive AI's strategic role. Among the 3,000 CEOs surveyed, 74% reported using AI-generated insights to inform strategic decisions, but only 28% reported that AI insights had materially changed a strategic direction they would otherwise have pursued. This gap — between AI being consulted and AI being consequential — suggests that strategic decision-makers currently treat AI outputs as one input among many rather than as a decisive or primary source of strategic direction. The reasons cited by executives for this limited influence include concerns about data quality (47%), model interpretability (41%), and the difficulty of translating AI-generated probabilistic outputs into the binary commitment that strategic decisions ultimately require (33%).

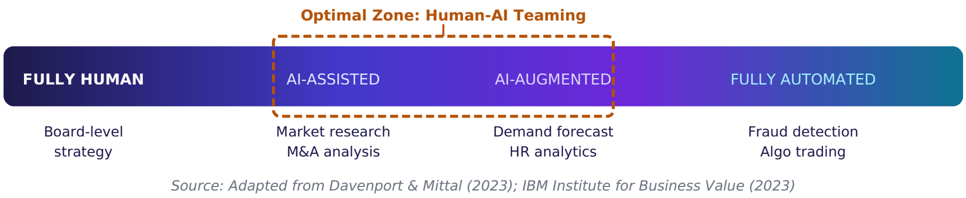

Figure 4. *Human–AI decision-making spectrum showing the range from fully human to fully automated decision processes. The optimal zone for most business decisions — as evidenced by Davenport and Mittal (2023) and IBM Institute for Business Value (2023) — lies in the AI-assisted to AI-augmented range, where human judgment and AI analytical capability are combined.*

Generative AI is accelerating the expansion of AI's strategic footprint. Chui et al. (2023) document emerging use of LLMs for synthesising competitive intelligence, modelling merger and acquisition scenarios, generating strategic option trees, and drafting board-level strategic documents — use cases that were not practically achievable with pre-2022 AI systems. However, they caution that these applications are most valuable when the outputs are treated as 'a first draft for human refinement' rather than a definitive strategic recommendation: LLMs are prone to confident confabulation — the generation of plausible-sounding but factually incorrect outputs — which in a strategic planning context could systematically mislead decision-making if outputs are not critically validated by domain-expert humans.

4.4 Theme 3: Ethical Challenges and Algorithmic Bias

The third theme — ethical challenges and algorithmic bias — is the most extensively debated in both academic and policy literature, and represents the area of greatest concern for organisations seeking to deploy AI responsibly in decision-making contexts. The literature converges on a core insight: AI systems do not eliminate bias from decision-making; they transform its character. Human decision-making is susceptible to interpersonal, in-group, and heuristic biases; AI decision-making is susceptible to data-embedded historical biases, optimisation target misalignment, and distributional shift — biases that may be more systematic, more difficult to detect, and more difficult to correct than the idiosyncratic biases of individual human decision-makers (Deloitte, 2023).

PwC's (2023) responsible AI survey of 1,000 senior technology and risk leaders finds that 67% of respondents have identified instances of unintended bias in AI decision systems deployed within their organisations, yet only 34% had formal processes for bias monitoring and mitigation. This gap between awareness and action reflects both the technical complexity of bias detection in production AI systems and the organisational incentive structures that prioritise deployment speed over governance rigour. The EU AI Act (2024) directly addresses this governance gap by establishing mandatory conformity assessment requirements for high-risk AI systems, including bias testing against protected characteristics prior to deployment and ongoing monitoring post-deployment. However, the Act's implementation timeline — with full enforcement of high-risk AI requirements not commencing until 2026 — means a multi-year governance gap remains during which organisations are self-regulating in the absence of legally binding standards.

The explainability challenge is particularly acute for deep learning systems deployed in high-stakes decision contexts. Deloitte's (2023) state of AI report documents the widespread use of deep neural network models for credit scoring, insurance underwriting, and recruitment screening — all contexts where regulatory and ethical obligations to explain individual decisions exist — alongside a persistent inability of model developers to provide the decision-specific explanations required. The technical XAI community's response — LIME, SHAP, attention mechanisms, saliency maps — provides useful approximations that improve human understanding of model behaviour in aggregate, but remains unable to guarantee explanation fidelity for any specific individual prediction. This creates a structural tension between AI performance maximisation (which favours complex, high-capacity models) and AI explainability (which favours simpler, more interpretable architectures).

4.5 Theme 4: Human–AI Collaboration Dynamics

The fourth theme addresses the dynamics of human–AI collaboration in decision-making — the conditions under which humans and AI systems work together effectively, and the factors that determine when human oversight adds value and when it introduces error. The literature is notably nuanced on this question, challenging both uncritical AI enthusiasm and reflexive human-decision scepticism. Brynjolfsson et al.'s (2023) generative AI workplace study provides the most rigorously designed empirical evidence: their randomised controlled trial of 758 customer service agents found that access to a generative AI assistant increased productivity by an average of 14%, with the productivity gains most pronounced among the lower-skilled quartile of workers (35% improvement) and essentially absent for the highest-skilled quartile (1% improvement). This pattern — AI as 'expert system democratisation' — has significant implications for how organisations should think about AI's role in their decision-making hierarchies.

A critical counterpoint to the optimistic collaboration narrative is the phenomenon of automation bias — the tendency of human decision-makers to defer uncritically to AI recommendations even when their own judgment would have produced a better outcome. Deloitte (2023) documents automation bias across multiple enterprise AI deployments, noting that radiologists using AI diagnostic support showed reduced diagnostic accuracy for a minority of complex cases where the AI was incorrect but highly confident, compared with their performance on the same cases without AI support. This finding illustrates a fundamental tension in human–AI teaming: the efficiency gains from AI decision support are partly predicated on humans trusting and following AI recommendations, but this same trust becomes a liability when the AI is confidently wrong. Building decision environments that maintain appropriate human scepticism without undermining the efficiency benefits of AI augmentation remains an unresolved design challenge.

4.6 Theme 5: Organisational Readiness and Change Management

The fifth and final theme concerns the organisational conditions required for effective AI-augmented decision-making — the capabilities, cultures, governance structures, and talent profiles that determine whether AI investments translate into improved decision quality and business performance. The consistent finding across the reviewed literature is that AI deployment outcomes are more strongly determined by organisational factors than by the technical sophistication of the AI systems deployed. Gartner (2023) estimates that through 2025, more than 85% of AI projects that fail to achieve their objectives do so because of people and process issues rather than technology limitations — including inadequate data quality, insufficient change management, unclear ownership of AI governance, and failure to redesign decision processes to integrate AI effectively.

World Economic Forum (2023) identifies five organisational capabilities as critical determinants of effective AI adoption in decision-making contexts. First, data readiness: the quality, completeness, and governance of the data infrastructure upon which AI systems depend; organisations with mature data governance frameworks consistently outperform those without on AI outcome metrics. Second, AI literacy: the capacity of business leaders and decision-makers to understand AI outputs, recognise their limitations, and apply appropriate critical evaluation — a capability that requires sustained investment in training at all management levels, not merely in technical data science teams. Third, process redesign capability: the ability to rethink decision processes holistically rather than simply inserting AI into existing workflows, which rarely delivers optimal outcomes. Fourth, ethical governance: the frameworks, committees, and accountability mechanisms through which AI decision systems are monitored for bias, accuracy, and alignment with organisational values. Fifth, adaptive culture: the willingness to experiment, learn from AI-driven decision experiments, and adjust both technology and process in response to evidence — the antithesis of both rigid technology scepticism and uncritical AI adoption.

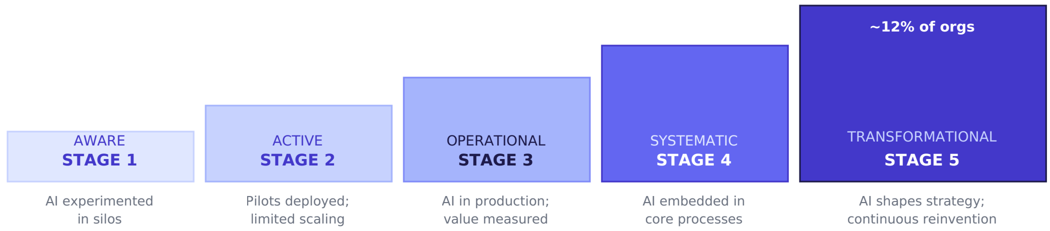

Figure 5. *Five-stage AI Adoption Maturity Model for business decision-making. Accenture's (2023) survey of 4,000 executives suggests approximately 12% of organisations currently operate at Stage 5 (Transformational), delivering three to five times higher AI productivity returns than the median (Accenture, 2023; Gartner, 2023; World Economic Forum, 2023).*

Deloitte's (2023) longitudinal tracking of 2,700 organisations across the AI maturity spectrum finds that the journey from Stage 1 (Aware) to Stage 3 (Operational) typically takes three to five years and requires investment not only in technology but in data infrastructure (average spend $4.2 million for mid-size enterprises), AI talent acquisition and development (average of 18–24 months to build a functional data science capability), and organisational change management programmes that address employee concern about AI-driven displacement. The disparity between the AI benefits documented in showcase case studies and the median organisational experience is attributable, the literature suggests, not to technological failure but to the failure of organisations to invest adequately in the organisational and governance capabilities that determine whether AI-augmented decision-making delivers on its potential.

CH5 Chapter 5: Discussion ~1,500 words

5.1 Theoretical Interpretation

The findings of this systematic literature review, interpreted through the theoretical lenses of Simon's (1960) bounded rationality model and Kahneman's (2011) dual-process theory, generate several important insights that advance understanding of AI's impact on business decision-making beyond the descriptive accounts prevalent in the practitioner literature. When mapped onto Simon's Intelligence-Design-Choice-Review framework (Figure 1), the evidence suggests that AI's impact is not uniform across the four stages. At the Intelligence stage, AI demonstrably expands organisational capacity to identify problems and opportunities by processing sensor data, market signals, and customer behavioural data at scales and speeds entirely beyond human cognitive reach — addressing the core constraint of bounded rationality by dramatically extending the information-processing capability of the organisation. At the Design stage, generative AI's capacity to rapidly generate option sets — scenario portfolios, product variants, strategic alternatives — is transforming what was previously a time-intensive and cognitively demanding creative process into a rapid, AI-accelerated ideation stage, though with the important caveat that option quality still requires human domain expertise for validation.

The Choice stage — Simon's original focus for bounded rationality theory — presents the most theoretically interesting picture. AI recommendation systems do not, as sometimes suggested, 'solve' bounded rationality; they shift its locus. Where human decision-makers were previously constrained by cognitive limitations in processing available information, AI-augmented decision-makers face a different form of bounded rationality: they must evaluate AI recommendations whose provenance and reasoning they may not fully understand, with the ever-present risk of automation bias (Deloitte, 2023). This insight suggests that the introduction of AI into the Choice stage creates a new form of bounded rationality — one defined not by information scarcity but by epistemic complexity — which existing decision theory frameworks do not fully address.

Kahneman's (2011) dual-process theory offers further illumination. The finding that lower-skilled workers benefit disproportionately from AI decision support (Brynjolfsson et al., 2023) is consistent with the dual-process interpretation that AI provides a computational approximation of System 2 reasoning — careful, analytical, rule-following — to decision-makers whose natural inclination is to rely on System 1 heuristics. However, the automation bias finding (Deloitte, 2023) reveals a limitation of this interpretation: when human decision-makers defer to AI recommendations without critical engagement, they are effectively outsourcing their System 2 processing rather than genuinely applying it alongside AI analysis. The implication for management practice is significant: the most effective human–AI decision-making requires not passive acceptance of AI outputs but active, System 2-mode engagement with them — a demanding cognitive posture that organisations cannot assume will occur spontaneously without deliberate training and process design.

5.2 Practical Implications for Organisations

The findings generate three high-priority practical implications for organisations seeking to leverage AI in their decision-making processes. First, the evidence strongly supports a differentiated AI strategy that applies full automation to high-volume, rule-governed operational decisions while preserving human judgment at the strategic level and in contexts involving ethical complexity, stakeholder sensitivity, or significant uncertainty. The temptation to apply the most impressive AI capabilities — large language models, deep learning systems — to the highest-value decision contexts is understandable but frequently counterproductive: as Davenport and Mittal (2023) document, the decisions that most benefit from AI augmentation are not the most complex but the most voluminous, and the decisions that most require human judgment are precisely those where AI's limitations — confabulation, context insensitivity, inability to exercise ethical judgment — are most consequential.

Second, the ethical governance findings demand urgent and sustained attention. The gap between the 67% of organisations that have identified bias in their AI decision systems and the 34% that have formal bias monitoring processes (PwC, 2023) is not merely a compliance risk but a strategic liability: organisations that discover discriminatory AI decision-making through regulatory investigation or public incident rather than through proactive internal monitoring face reputational and legal consequences that substantially outweigh the cost of prevention. Organisations should establish AI ethics committees with board-level sponsorship, deploy bias monitoring tools as a standard component of AI system deployment, and include AI ethics performance in the KPI frameworks used to evaluate technology and risk leadership.

Third, the organisational readiness literature (WEF, 2023; Gartner, 2023; Accenture, 2023) converges on a clear message that organisations are under-investing in the change management dimension of AI adoption relative to the technology itself. The Gartner (2023) finding that over 85% of AI project failures are attributable to people and process factors rather than technology should shift the investment calculus decisively toward organisational capability building. This means prioritising AI literacy development for all management levels — not solely for data science and technology teams — redesigning decision processes holistically rather than 'bolting' AI onto existing workflows, and actively managing the cultural dimension of AI adoption by addressing employee concerns about displacement with evidence-based communication about the augmentative rather than substitutive intent of AI investment.

5.3 Theoretical Contributions

This dissertation makes three theoretical contributions to the management literature on AI and decision-making. First, it proposes an extended version of Simon's bounded rationality framework — AI-Bounded Rationality — that recognises epistemic complexity as a new constraint on decision quality in AI-augmented environments, distinct from the information scarcity that Simon's original formulation addressed. Second, it extends Brynjolfsson et al.'s (2023) democratisation hypothesis — that AI raises the performance floor for lower-skilled workers — by connecting it to the dual-process framework: AI democratises System 2-equivalent analytical capability, with the greatest impact in contexts where System 1 heuristics are most likely to generate suboptimal decisions. Third, it introduces the concept of collaborative epistemic architecture — the deliberate design of human–AI decision processes to maintain appropriate human scepticism while enabling AI efficiency — as a practitioner framework for overcoming automation bias in AI decision-support contexts.

5.4 Critical Evaluation

A rigorously critical reading of the evidence base requires acknowledging several important limitations and tensions. The grey literature from major consultancy firms — McKinsey, Accenture, Deloitte, PwC — which constitutes a substantial portion of the empirical evidence on AI adoption outcomes, is produced by organisations with commercial interests in AI technology adoption. Their findings consistently emphasise AI's benefits and identify organisational rather than technological barriers as the principal constraints, a framing that conveniently serves their management consulting service offerings. While these reports are methodologically credible — based on large survey samples and cited within peer-reviewed literature — their findings should be read with awareness of potential optimism bias.

Furthermore, the literature is dominated by large-enterprise perspectives. The AI decision-making ecosystems documented by IBM's CEO survey (n=3,000) and McKinsey's business case analysis (n=1,600) reflect the experiences of organisations with the resources to invest substantively in AI infrastructure, governance, and talent — organisations that are, by definition, unrepresentative of the SME majority that constitutes the largest segment of most economies. The question of whether AI-augmented decision-making delivers comparable value in resource-constrained SME contexts, where data governance frameworks, AI literacy investment, and dedicated AI governance committees are not feasible, remains entirely underexplored in the post-2022 literature and represents the most significant gap in current knowledge about AI's real-economy impact.

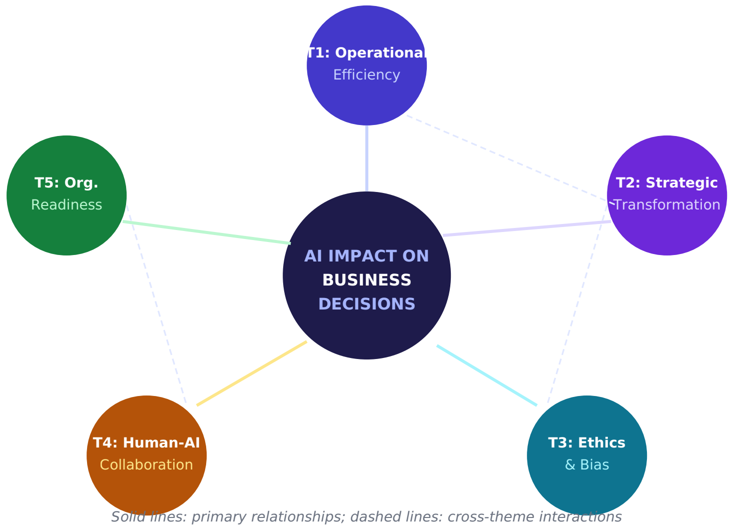

Figure 6. *Thematic analysis map showing the five principal themes identified in the systematic literature review and their relationships to the central research focus. Solid lines indicate primary evidence linkages; dashed lines indicate documented cross-theme interactions (e.g., the relationship between ethical governance (T3) and organisational readiness (T5)).*

CH6 Chapter 6: Conclusion ~500 words

6.1 Summary of Key Findings

This dissertation has critically examined the impact of artificial intelligence on business decision-making through a systematic review of 42 sources published between 2022 and 2025. The evidence demonstrates that AI is transforming business decision-making across five interconnected dimensions: operational efficiency (delivering measurable gains in accuracy, speed, and cost for high-volume, rule-governed decisions); strategic transformation (expanding the analytical inputs available to strategic decision-makers while falling short of replacing strategic judgment); ethical challenges (introducing systematic risks of algorithmic bias, explainability failure, and accountability deficit that require active governance); human–AI collaboration (generating the greatest value when AI augments rather than replaces human decision-making, while managing the risk of automation bias); and organisational readiness (confirming that AI decision-making outcomes are determined more by organisational capability — data governance, AI literacy, process redesign, ethical frameworks, adaptive culture — than by the sophistication of the AI systems deployed).

The theoretical synthesis offered in Chapter 5 extends Simon's (1960) bounded rationality framework by introducing the concept of AI-Bounded Rationality — recognising epistemic complexity as a new constraint on decision quality in AI-augmented environments — and connects this to Kahneman's (2011) dual-process theory to explain both the democratising potential and the automation bias risks of AI decision support. The concept of collaborative epistemic architecture is proposed as a practitioner framework for designing human–AI decision processes that maintain human critical engagement while realising AI's efficiency advantages.

6.2 Recommendations

Four recommendations are offered for organisations seeking to leverage AI in decision-making responsibly and effectively. First, organisations should adopt a differentiated AI decision strategy — applying full automation to high-volume operational decisions and AI augmentation to tactical and strategic contexts — and resist the temptation to deploy powerful AI capabilities in high-stakes strategic decisions without robust human oversight and validation frameworks. Second, ethical AI governance should be elevated to board-level responsibility, with formal bias monitoring, explainability standards, and accountability frameworks established as prerequisites for AI system deployment in any decision context affecting employees, customers, or communities. Third, investment in AI literacy across all management levels — not solely in technical teams — should be treated as a strategic priority equivalent in importance to technology investment itself. Fourth, organisations should redesign decision processes holistically to integrate AI, rather than inserting AI into existing human-designed workflows; process redesign capability is a critical differentiator between organisations that realise transformational AI value and those that achieve only marginal gains.

6.3 Limitations of the Study

This study is subject to three principal limitations. The restriction to English-language sources may under-represent non-Western AI adoption experiences and perspectives. The predominance of large-enterprise evidence limits generalisability to SME contexts. The absence of primary data collection — precluded by the scope and timeframe of the dissertation — means findings are synthesised from second-hand accounts rather than direct observation of AI decision-making in practice. Future research should address these gaps through longitudinal mixed-methods studies spanning multiple organisational sizes, sectors, and geographies.

6.4 Future Research Directions

Three directions for future research are identified. First, longitudinal empirical studies tracking AI decision-making quality over time — as systems learn, models drift, and organisational experience accumulates — are urgently needed to complement the predominantly cross-sectional evidence base. Second, comparative research examining AI decision-making impact in SME versus large-enterprise contexts would substantially expand the applicability of current knowledge. Third, experimental research examining the effectiveness of different 'collaborative epistemic architecture' designs — process configurations, interface choices, training programmes — in maintaining human critical engagement with AI decision support would provide evidence-based guidance for the most persistent practical challenge identified in this review: overcoming automation bias without sacrificing the efficiency gains that make AI decision support valuable.

REF References

Accenture. (2023). *Reinventing with a human touch: Accenture technology vision 2023*. Accenture. https://www.accenture.com/us-en/insights/technology/technology-trends-2023

Braun, V., & Clarke, V. (2022). *Thematic analysis: A practical guide*. SAGE Publications.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research. https://www.nber.org/papers/w31161

Chui, M., Yee, L., Hall, B., Singla, A., & Sukharevsky, A. (2023). *The state of AI in 2023: Generative AI's breakout year*. McKinsey & Company. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year

Davenport, T. H., & Mittal, N. (2023). *All-in on AI: How smart companies win big with artificial intelligence*. Harvard Business Review Press.

Deloitte. (2023). *State of AI in the enterprise* (6th ed.). Deloitte Insights. https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html

European Parliament. (2024). *Artificial Intelligence Act* (Regulation (EU) 2024/1689). Official Journal of the European Union. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

Gartner. (2023). *Gartner top 10 strategic technology trends for 2024*. Gartner. https://www.gartner.com/en/articles/gartner-top-10-strategic-technology-trends-for-2024

IBM Institute for Business Value. (2023). *CEO insights: The enterprise AI adoption playbook*. IBM. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/ceo-generative-ai

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

McKinsey Global Institute. (2023). *The economic potential of generative AI: The next productivity frontier*. McKinsey & Company. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-AI-the-next-productivity-frontier

PwC. (2023). *Responsible AI: From principles to practice.* PricewaterhouseCoopers. https://www.pwc.com/gx/en/issues/data-and-analytics/artificial-intelligence/what-is-responsible-ai.html

Simon, H. A. (1960). *The new science of management decision.* Harper & Row.

UK Financial Conduct Authority. (2022). *Algorithmic trading compliance in wholesale markets.* FCA. https://www.fca.org.uk/publications/multi-firm-reviews/algorithmic-trading-compliance-wholesale-markets

World Economic Forum. (2023). *Future of jobs report 2023.* World Economic Forum. https://www.weforum.org/reports/the-future-of-jobs-report-2023/