Stockholm School of Economics
1351 Master's thesis in Business Management
Academic year 2025–2026

The Promised Goldrush

A Multiple-Case Study of How AI Agents can Create Value Across the Innovation Funnel

Leo Hall (42807) and Jakob Vaghult (42551)

Supervisor: Korinzia Toniolo
Date examined: 17th of December - 2025
Discussant(s): Felix Fröndhoff & Bror Nordström
Examiners: Emma Bell & Örjan Sjöberg

Abstract

Autonomous AI agents are creating increasing value in organizations, yet this remains uneven with few organizations actually unlocking this value. The technology presents a particularly promising value-add in innovation processes, yet research offers limited understanding of how these technologies can create value and what conditions shape their realization. Addressing this gap, this thesis investigates how agentic AI contributes across the innovation funnel and why similar technologies yield divergent outcomes across firms. Drawing on affordance theory and seven case studies spanning industry-native, digital-native, and AI-native organizations, the analysis identifies five affordance potentials through which agents can generate innovation value and a corresponding set of technical, organizational, and domain-level actualization conditions that govern whether this value materializes. Mapping these dynamics across the innovation funnel shows that agentic value appears stage-contingent: most affordance potential lies upstream, with non-digital downstream implementation phases presenting more limited opportunities for contemporary AI agents. Integrating these findings, the thesis develops the Agentic Value Ladder, a mid-range model explaining how organizations progress from low-order efficiency improvements to higher order affordance potentials creating new forms of capabilities from AI agents. Together, the findings extend AI Agent and affordance theory literature with a theoretically grounded account of how AI agents create value across the innovation funnel as well as how conditions constrain and enable the actualization of this value.

Keywords: AI agents, Agentic AI, Innovation Management, Innovation Funnel, Affordance

Acknowledgements

We would like to express our sincere gratitude to our supervisor Korinzia Toniolo for her guidance, thoughtful feedback, and intellectual support throughout this thesis. Her expertise in AI and technological innovation research substantially strengthened our analytical approach and contributed meaningfully to the development of our theoretical insights.

We also extend our thanks to all company respondents who generously shared their time, perspectives, and experiences. Their openness and depth of reflection made this study possible and provided the empirical foundations for understanding the emerging role of AI agents in innovation processes.

We are additionally grateful for the constructive discussions with our peers at Stockholm School of Economics and with supporting technical experts from institutions and firms, whose perspectives on AI systems helped us refine and stress-test elements of our analysis.

Table of Contents

Abstract 2
Acknowledgements 3
Definitions 6
1. Introduction 7
    1.1 Background and Motivation 7
    1.2 Research Gap and Problem Definition 8
    1.3 Purpose and Research Questions 9
    1.4 Expected Contribution 10
2. AI agents 11
    2.1 Generations of Artificial Intelligence 11
    2.2 Defining and Structuring AI agents 12
3. Literature review 14
    3.1 The Innovation Funnel 14
    3.2 Affordance Theory - Understanding AI agents' Potential and Constraints 18
    3.3 Research Gap 20
4. Methodology 21
    4.1 Research Design 21
    4.2 Data Collection 22
    4.3 Data Coding & Analysis 24
    4.4 Quality Considerations 25
    4.5 Ethical Considerations 26
5. Findings 27
    5.1 Case Overviews 27
    5.2 Cross-Case Empirical Findings 32
    5.3 Cross-Case Comparison - Affordances across the Innovation Funnel 33
6. Analysis 42
    6.1 Emergent Theoretical Constructs 42
    6.2 Affordance Potential and Actualization Conditions across the Innovation Funnel 45
    6.3 The Agentic Value Ladder 49
    6.4 Mechanisms of the Agentic Value Ladder 51
7. Discussion 53
    7.1 Conclusion 53
    7.2 Theoretical Contributions 54
    7.3 Managerial Implications 56
    7.4 Limitations 57
    7.5 Future Research 58
8. Bibliography 59
Appendices 66
    Appendix A - Semi-Structured Interview Protocol 66
    Appendix B - Respondents Summary 67

Appendix C - Case Category Patterns and Quotes..... 69
Appendix D - AI Disclosure..... 72

Definitions

Large language models	Large language models (LLMs) are advanced neural networks trained on a vast collection of text to recognize patterns, generate language, and perform complex reasoning tasks (OpenAI, 2022).
Agentic AI	The most recent paradigm of AI systems, combining foundation models with planning, memory, tool use, and environment interaction to pursue goals with limited supervision (IBM, 2025; Nisa et al., 2025). It denotes the broader paradigm of autonomous reasoning and action.
AI agent	The concrete instantiations of the agentic AI paradigm within organizational settings. Defined as an automated entity that senses and responds to its environment and takes actions to achieve its goals (ISO, 2025: Nisa et al., 2025). These systems integrate a cognitive core (LLM), a tool interface, and an execution layer (IBM, 2025; Microsoft, 2024).
Innovation Funnel	A collectively exhaustive representation of how innovation progresses from broad exploration to validated solutions through successive stages that reduce uncertainty (Dunphy et al., 1996). The funnel is commonly represented through frameworks such as the Double Diamond (Design Council, 2005).
IT Architecture	IT architecture is the blueprint or design of how IT elements (infrastructure, applications, data, services) are arranged, interconnected, and managed, so they serve the enterprise's strategic and operational goals (Foong et al., 2013).
Affordance	Affordance refers to the potential for purposeful action that emerges from the relationship between a technology and the actors who use it. It is not an inherent property of the artefact nor solely a user attribute, but a relational possibility shaped by the material features of the technology and the social, organizational, and interpretive structures in which it is embedded (Hutchby, 2001; Markus & Silver, 2008).

1. Introduction

1.1 Background and Motivation

Innovation is widely recognized as a central engine of organizational renewal, competitiveness, and long-term economic growth (Schumpeter, 1934; Tidd & Bessant, 2018). Across industries, managers navigating technological change consistently confront two foundational questions: (1) what value-creation opportunities a new technology affords, and (2) what organizational conditions must be in place to capture that value. These questions reflect established findings that technology adoption alone does not generate advantage; value arises only when technologies are embedded within complementary organizational structures and capabilities (Brynjolfsson & Hitt, 2000; Teece, 2007). They become especially acute in the case of emerging general-purpose technologies, where potential is high but pathways to value remain uncertain (Bresnahan & Trajtenberg, 1995).

Artificial intelligence (AI) exemplifies such a technology. Over the past decade, AI has expanded from statistical prediction toward creativity, design, reasoning, and decision-making (Brynjolfsson & Mitchell, 2017; Mariani & Dwivedi, 2024). This has been enabled through Large language models (LLMs), the core technology of AI applications like ChatGPT with capabilities to reason and produce content (OpenAI, 2022; Mariani & Dwivedi, 2024; Teece, 2007). Building on these advances, developers are now constructing AI agents - systems that combine LLM-based reasoning with tool orchestration to autonomously execute multi-step workflows (Park et al., 2023; IBM, 2025). This architecture marks a qualitative departure from earlier rule-based automation which, without the reasoning capacities of LLMs, was limited to repetitive, deterministic tasks (Aguirre & Rodriguez, 2017).

The economic significance of this shift is substantial. Investors increasingly frame agentic systems as the next foundational technology, leading to a sharp rise in venture investment in agent-based startups in recent years (Dealroom, 2025; Sequoia Capital, 2025; Tech Monitor, 2025). Policymakers similarly emphasise AI systems as strategically important for innovation capacity and economic competitiveness (European Commission, 2024; White House OSTP, 2023).

Beyond this, AI agents introduce a particularly compelling opportunity within innovation processes. Despite their strategic importance, innovation activities are frequently deprioritised due to the persistent tension between long-term exploratory work and short-term operational demands - an imbalance captured in the classic exploration–exploitation Innovator's Dilemma

(March, 1991). In this context, agents' ability to autonomously scan signals, generate concepts, test prototypes, and iterate across multi-step workflows offers a novel way to sustain exploratory activity without continuously diverting human attention from ongoing operations (Park et al., 2023; Tekic & Füller, 2023). While automation is not new, innovation work, which involves high levels of reasoning, ambiguity, and non-repetitive problem solving, had limited use of the previous deterministic rule-based systems. The reasoning capabilities enabled by contemporary LLM-based agents therefore represent a meaningful shift, with early evidence suggesting that such agents can accelerate prototyping, broaden ideation, and enhance development feedback loops (Crossan & Apaydin, 2010, Händler, 2023; Humberd & Latham, 2025; Hobday, 2005).

Taken together, these developments suggest that AI agents may hold distinctive value for innovation processes - yet clarity in how to enable this potential is lacking. While AI-linked firms already account for a disproportionate share of global equity growth (Financial Times, 2025; The Economist, 2024), measured productivity effects remain modest and uneven, heavily dependent on firms' digital maturity (Brynjolfsson et al., 2025). This mismatch between market expectations and realised organizational outcomes highlights the importance of understanding how and under what conditions AI agents can generate value in innovation processes.

1.2 Research Gap and Problem Definition

Very little work examines AI as an autonomous actor capable of operating independently across stages of the innovation process. Existing studies almost exclusively conceptualise AI as an assistive tool for discrete, task-level augmentation - for example within data-driven problem discovery (Mariani & Dwivedi, 2024; Truong & Papagiannidis, 2022), predictive evaluation during problem definition (Oberländer et al., 2021; Shrestha et al., 2019), or generative concept development in prototyping and design (Cimino et al., 2024). Systematic evidence on agentic, orchestrating, or multi-step AI systems remains largely absent, a gap highlighted by recent work calling for research on autonomous AI within innovation processes (Füller et al., 2024; Gama & Magistretti, 2023; Tekic & Füller, 2023). At the same time, value-creation research emphasises that technological autonomy does not inherently yield organizational benefit; it becomes valuable only when aligned with governance structures, data infrastructures, and routines (Lepak et al., 2007).

The resulting gap concerns both potential and realisation: we lack empirical insight into where in innovation processes agentic AI can create value, and we lack theoretical understanding of the organizational conditions that shape whether that value is captured. The accelerating investment and strategic attention surrounding agentic systems further amplify the importance of

addressing this gap. To analyse this dual challenge, a theoretical lens is needed that accounts both for the action possibilities introduced by a technology and for the organizational conditions that shape their enactment. Affordance–actualization theory offers such an initial foundation (Gibson, 1977; Hutchby, 2001; Markus & Silver, 2008; Strong et al., 2009; Volkoff & Strong, 2017). Worthy of note, is that existing theory remains limited in important ways - particularly its predominant focus on micro-level contexts (Fayard & Weeks, 2014; Keller et al., 2019) and their limited treatment of organizational preconditions (Besson & Rowe, 2012). Nonetheless, its emphasis on technological possibilities and constraints renders it a suitable foundation.

The gap in the literature is not bound to one specific part of innovation processes, highlighting the need for a comprehensive review across the innovation funnel. The innovation funnel represents a collectively exhaustive representation of how innovation progresses from broad exploration to validated solutions through successive stages that reduce uncertainty (Dunphy et al., 1996). Given the nascency of autonomous AI in innovation, such a holistic, funnel-spanning perspective is necessary to establish an initial, system-level understanding of how agentic capabilities interact with the dynamics of innovation work, thereby clarifying where deeper, stage-specific research should subsequently focus.

1.3 Purpose and Research Questions

Against this backdrop, the purpose of this thesis is to clarify where AI agents can create value within the innovation funnel and to identify the organizational conditions that shape whether that value can be realised in practice. Building on this purpose, two key research questions are presented:

RQ1: *How can AI agents create value across the innovation funnel?*

RQ2: *How do organizational conditions enable or constrain the value potential of AI agents across the innovation funnel?*

To address these questions, we conduct a multiple-case study of Nordic organizations across three archetypes - industry-native, digital-native, and AI-native firms - each corresponding to a distinct technological basis for their core business and offerings: physical (non-digital), digital, and AI agent–based technologies. Semi-structured interviews with managers and technical experts enable comparison across varying levels of digital maturity, revealing how socio-technical environments condition the realization of agentic value (Gama & Magistretti,

2023; Hoessler et al., 2024; Oberländer et al., 2021). Through iterative within- and cross-case analysis, we develop an integrative framework linking the stages of the innovation funnel with the value creation opportunities AI agents present and their corresponding enabling and constraining conditions.

1.4 Expected Contribution

This thesis contributes to emerging research on agentic AI in innovation by advancing understanding in three ways. Conceptually, it extends affordance–actualization perspectives to autonomous AI systems, identifying stage-specific innovation affordances and the organizational boundary conditions that shape their enactment. Empirically, it enriches the limited evidence base through a comparative multiple-case study of industry-native, digital-native, and AI-native firms, illustrating how agentic capabilities manifest differently across socio-technical contexts. Practically, it can offer knowledge supporting innovation leaders in assessing how agents can generate value, which organizational enablers are required, and why deployments may underperform. Together, these contributions provide a foundation for more context-sensitive integration of AI agents into innovation processes.

2. AI agents

This section establishes the conceptual foundation for understanding AI agents as a distinct technological paradigm. It traces the evolution of AI from rule-based systems to predictive, generative and agentic architectures (2.1), highlighting the shift from structured logic to adaptive reasoning and autonomous action. It then defines AI agents and outlines the architectural layers that enable perception, planning and execution in organizational settings (2.2).

2.1 Generations of Artificial Intelligence

The evolution of Artificial Intelligence can be conceptually framed through four paradigms: symbolic, sub-symbolic (predictive), generative, and agentic. Together, they trace a trajectory from human-encoded logic toward autonomous, goal-directed systems capable of perception, reasoning, and action.

First, Symbolic AI represents knowledge through human-authored rules and formal logic structures, allowing systems to make decisions by applying explicit "if-then" logic written by domain experts (Alavi & Leidner, 2001; Russell & Norvig, 2021). These systems delivered transparent reasoning, but they were brittle because they could only operate within the predefined rule base and could not adapt when real-world cases deviated from the encoded conditions.

Second, sub-symbolic or predictive AI addressed these limitations by enabling algorithms to learn statistical patterns directly from data, rather than relying on fixed human-authored rules. In these systems, models infer relationships - such as which borrower characteristics predict default risk - by analyzing large numbers of past examples (Soori et al., 2023). Neural networks and deep-learning architectures further expanded machine cognition from explicit rules to flexible, data-driven inference (Soori et al., 2023).

Third, Generative AI extended these capabilities by allowing models not only to classify or predict, but to produce new artefacts such as text, imagery, or code. Transformer-based models learned complex data distributions and enabled a wide range of content creation (IBM, 2025). These developments provided the foundation on which more autonomous systems could be built.

Fourth, agentic AI represents the most recent paradigm. These systems combine foundation models with planning, memory, tool use, and environment interaction to pursue goals with limited supervision (IBM, 2025). Following Nisa et al. (2025), this thesis uses agentic AI to denote the broader paradigm of autonomous reasoning and action, and AI agents as the

concrete instantiations of that paradigm within organizational settings. Table 1 summarises these four paradigms and clarifies how AI agents operationalise the agentic AI paradigm.

Table 1 - Generations of AI Paradigms

Paradigm	**Core Capability**	**Typical Examples**
Symbolic AI	Executes rule-based reasoning on explicit knowledge representations.	Early knowledge management systems (Russell & Norvig, 2021).
Sub-symbolic / Predictive AI	Learn statistical patterns from data to predict or classify outcomes.	Credit-risk models, image recognition, recommendation engines (Jordan & Mitchell, 2015)
Generative AI	Produces novel content (text, image, code) by modelling data distributions and context.	ChatGPT & GitHub Copilot (Eloundou et al., 2023).
Agentic AI	Perceives, plans, and acts autonomously toward goals, integrating reasoning and execution.	A Market Analysis Agent scraping the web and CRM systems to generate reports or general purpose Microsoft Co-Pilot Agents (Microsoft, 2024)

This evolution reflects a steady increase in autonomy and integration. Whereas symbolic systems reasoned within human-specified rule spaces, predictive systems learned statistical regularities, and generative systems created new artefacts from those regularities, agentic systems act: they autonomously pursue objectives, coordinate subsystems, and close the loop between insight and execution. The following sections build on this conceptual foundation to define what constitutes AI agents and how they differ from earlier automation technologies.

2.2 Defining and Structuring AI agents

The ISO (2025) defines an AI agent as an automated entity that senses and responds to its environment and takes actions to achieve its goals. Earlier rule-based agents were reactive and brittle. Modern agents rely on large language models (LLMs) that support context-sensitive

reasoning and multi-step planning, enabling autonomous perception, reasoning, and action (IBM, 2025).

Technically, AI agents integrate three layers that differentiate them from deterministic automation like Robotic Process Automation. First, the cognitive core (typically an LLM) interprets context and produces reasoned outputs (Microsoft, 2024). Second, the tool and context interface links the model to external systems through Application Programming Interface (APIs), Retrieval-Augmented Generation (RAG) pipelines, or emerging agent orchestration standards such as Model Context Protocol (MCP) (Anthropic, 2024; Amazon Web Services, 2025). Third, the execution and feedback layer monitors outcomes, invokes tools, and adapts future actions (Microsoft, 2024). This architecture shifts agents from one-shot output generation to iterative, goal-directed workflow execution. In this perspective, AI agents, with their probabilistic reasoning and adaptive tool use, differ from previous automation technologies that follow fixed rules (Aguirre & Rodriguez, 2017). This enhanced flexibility is particularly relevant in innovation work, where uncertainty, exploration, and judgment dominate.

Agents may operate as single systems or in multi-agent configurations. While multi-agent setups can expand capability, recent scholarship notes that they are not inherently "more agentic" (Nisa et al., 2025). The degree of agency depends on autonomy and coherence in goal pursuit, not on the number of agents. A well-integrated single agent can outperform a multi-agent system requiring extensive human intervention. This aligns with the ISO view of agency as a spectrum of autonomy rather than an architectural label (ISO, 2025).

3. Literature review

This chapter reviews the literature that underpins the thesis. It begins by outlining established models of the innovation process and the funnel structure used to organize our analysis (3.1). It then examines how existing research conceptualises the role of AI in innovation and how these insights connect with broader affordance–actualization perspectives (3.2). Finally, it synthesizes the unresolved issues across these streams to articulate the research gap that motivates the study (3.3).

3.1 The Innovation Funnel

3.1.1 Foundations and Logic of the Innovation Funnel

The innovation-management literature offers numerous frameworks describing how ideas evolve from exploration to implementation, yet most converge on a socio-technical logic of iterative divergence and convergence through which uncertainty is progressively transformed into actionable outcomes (Hobday, 2005; Tidd & Bessant, 2018). Early linear models such as Schumpeter's (1934) creative destruction and Rogers' (1962) diffusion framework conceptualised innovation as a sequential process, but subsequent empirical research demonstrated that organizational innovation rarely unfolds in a straight line.

Building on this shift toward non-linear models, the Design Council introduced the Double Diamond model (2005). Derived from ethnographic studies in design-intensive firms such as IDEO and Philips, it distilled innovation into two successive divergence–convergence cycles, capturing the oscillation between exploratory search and evaluative refinement. Its emphasis on problem reframing, iterative testing, and socio-technical coordination has contributed to its widespread adoption across industries and scholarship.

In line with the Double Diamond model, this thesis conceptualises innovation as a four-stage funnel comprising Discover, Define, Develop and Deliver. Discover covers broad exploration, signal detection and opportunity identification. Define focuses on convergence through prioritisation and problem framing. Develop involves iterative refinement, prototyping and validation. Deliver concerns implementation, scaling and organizational diffusion. Framing the process in this way offers a clear analytical structure for examining how digital technologies reshape innovation work across stages. Although innovation rarely proceeds in such clearly defined stages in practice, this level of abstraction enables meaningful comparison across our cases. The funnel provides an analytically tractable structure that captures core

divergence–convergence dynamics while allowing differences in socio-technical context to be examined systematically.

3.1.2 AI Across the Stages of the Innovation Funnel

While research on AI in innovation management has expanded rapidly, studies remain distributed across specific tools, tasks, and phases of the innovation process. To clarify how different AI capabilities influence innovation work, this section synthesizes empirical evidence across the four stages of the innovation funnel. Existing research shows a progression from early analytic augmentation toward generative creativity and, more recently, preliminary explorations of autonomous behaviours. Examining these developments stage-by-stage highlights both the expanding scope of AI support and the fragmented nature of current findings, providing the basis for identifying cross-stage gaps in Section 3.1.2.5.

3.1.2.1 AI in the Discover Stage

The Discover stage concerns broad opportunity exploration, weak-signal detection, and the generation of initial concepts. Early work on AI in innovation positioned machine learning and data-driven analytics as tools for strengthening exploratory search. Research shows that these systems can broaden firms' sensing capabilities by helping them detect emerging trends, user needs, and market signals more quickly and at larger scale than traditional approaches (Du et al., 2020; von Krogh, 2018). In parallel, practical applications documented by Mariani & Dwivedi (2024) and Truong and Papagiannidis (2022) demonstrate how AI-supported market analysis and idea filtering improved ideation efficiency and decision confidence, though organizational inertia and limited absorptive capacity often constrained adoption.

The emergence of generative AI broadened the Discover-stage search space even further. Generative models have been shown to accelerate concept generation, synthesize diverse inputs, and encourage more radical exploration by reducing the cost of producing early-stage ideas (Cimino et al., 2024; Gama & Magistretti, 2023; Mariani & Dwivedi, 2024). These systems enable rapid, parallel experimentation across thematic directions that would be difficult or costly for human teams to explore manually. However, most studies focus on tool-level augmentation rather than how generative systems integrate into established discovery routines or organizational search processes.

3.1.2.2 AI in the Define Stage

In the Define stage, innovation activities focus on structuring, evaluating, and prioritising the ideas surfaced during Discover. Predictive and classification models aid this convergence by transforming large volumes of user, market, and technical data into forecasts, relevance scores, and comparative indicators that help teams screen options and identify high-potential directions (Shrestha et al., 2019; Truong and Papagiannidis, 2022). Experimental work further shows that AI-based advisory systems can increase managers' confidence in selection decisions, though such systems may also induce overreliance due to the perceived analytical rigour of algorithmic outputs (Keding and Meissner, 2021). Generative AI extended these capabilities by consolidating heterogeneous stakeholder inputs, extracting latent themes, and structuring qualitative feedback to clarify differences between competing ideas (Gilardi, Alizadeh and Kubli, 2023; Şen, Şen and Şahin, 2023). Generative AI also supports iterative refinement of early concepts - helping articulate value propositions, elaborate design directions, and stress-test hypotheses - while final selection remains a human responsibility (Chen et al., 2025).

3.1.2.3 AI in the Develop Stage

The Develop stage involves iterative prototyping, design exploration, and experimentation. This stage has attracted substantial attention in recent research due to the alignment between generative capabilities and design-intensive activities. Generative design tools, simulation platforms, and multimodal generative models have been shown to accelerate variant generation, prototype iteration, and technical exploration (Cimino et al., 2024). These systems reduce the cost of exploration by enabling rapid experimentation with alternative forms, functions, and system architectures. LLM-based coding assistants increasingly facilitate rapid prototyping of digital products, lowering iteration friction and shortening feedback loops (Eloundou et al., 2023). At the same time, studies highlight integration challenges: generative design outputs are not always interoperable with downstream engineering tools, governance frameworks struggle to keep pace with generative experimentation, and teams face new issues related to versioning, authorship, and model bias (Roberts et al., 2024; Tekic & Füller, 2023). As in earlier stages, most research examines isolated tools rather than coordinated workflows.

3.1.2.4 AI in the Deliver Stage

Compared to upstream stages, research on AI in the Deliver stage - implementation, scaling, and organizational diffusion - remains limited. Existing studies describe data-driven decision-support applications in operational contexts, though these remain narrower in scope

than earlier-stage uses of AI (Haefner et al., 2021). A growing body of implementation research shows that organizations face substantial socio-technical barriers when integrating AI into production environments, including legacy system incompatibilities, workflow misalignment, and governance requirements that slow or constrain deployment (Jöhnk et al., 2021). Evidence from industrial and organizational contexts also indicates that AI-enabled analytical improvements do not automatically translate into realised outcomes, as value capture is often hindered by bottlenecks in cross-functional coordination, change management, and organizational readiness (Bailey & Barley, 2020; Enholm et al., 2022). Taken together, this suggests that the Deliver stage remains the least studied and least transformed part of the innovation funnel in existing AI-innovation literature (Haefner et al., 2021).

3.1.2.5 Cross-Stage Synthesis and Emerging Gaps

Across the four stages, the empirical literature presents a fragmented and task-specific view of AI's role in innovation. Discover-stage studies emphasise data-driven sensing and generative ideation; Define-stage research highlights predictive evaluation and conceptual synthesis; Develop-stage work focuses on generative design, coding assistance, and simulation; and Deliver-stage studies address forecasting and operational integration. While these findings show that AI can augment specific activities, they rarely examine how AI shapes interdependencies, feedback loops, or coordination across stages of the innovation process.

Two overarching gaps emerge. First, existing studies overwhelmingly investigate assistive AI- analytics, prediction, and generative content creation- rather than systems capable of autonomously coordinating multi-step innovation tasks. Despite early explorations of AI agents in manufacturing and ideation contexts, there is no systematic evidence on how increasing autonomy or tool orchestration affects innovation workflows as a whole. Second, current research pays limited attention to the organizational conditions that shape whether AI capabilities translate into realised outcomes. Factors such as governance routines, data interoperability, workflow design, and cross-functional coordination repeatedly appear as moderators of impact, yet remain theoretically underdeveloped.

Taken together, the literature lacks an integrated understanding of how AI agents can create value across the innovation funnel and what conditions enable these capabilities to be actualised. This gap also underscores the importance of theory capable of explaining both action potentials and the constraints that shape their realisation; a need to which affordance theory offers an initial but valuable conceptual foundation

3.2 Affordance Theory - Understanding AI agents' Potential and Constraints

3.2.1 Affordances as possibilities of technologies

agentic AI introduces a distinct class of technologies where performance relies less on real-time user input and more on the alignment between the system's autonomy and the organization's existing structures. This demands a theoretical approach that treats technology not as a fixed set of functionalities, but as a bundle of action possibilities dependent on the structural conditions surrounding them. Affordance theory provides this perspective. Originating in ecological psychology as the study of action possibilities (Gibson, 1977), and later extended in information systems research to examine how technologies enable and constrain goal-directed action (Hutchby, 2001; Markus & Silver, 2008), the affordance lens can support in explaining not only what agentic systems could do in principle, but how the interaction between social and technical structures shape what they can achieve in practice.

An affordance is the potential for goal-oriented action that emerges from the relationship between a user and a technology (Leonardi, 2011). These arise through interaction between technologies and organizational practice rather than residing in either alone (Orlikowski & Scott, 2008). Leonardi (2011) demonstrates how technologies and routines become "imbricated" as actors adjust practices in response to technological possibilities, and prior work consistently shows that identical systems may afford different actions across organizations depending on local routines, norms, and interpretive frames (Volkoff & Strong, 2017). This relational view is particularly relevant for innovation work, which involves iterative sensemaking, distributed responsibilities, and cross-functional coordination (Crossan & Apaydin, 2010; Hussmo, 2022).

3.2.2 Affordance actualization

Building on this foundation, Strong et al. (2009) describe actualization as the situated sequence of actions through which actors recognize an affordance, experiment with it, and work to realise it in practice. Actualization is therefore an enacted process rather than a technological property. Prior research acknowledges that various organizational and technical factors - such as data quality (Chen et al., 2008), system interoperability (Vernadat, 2009), governance structures (Weill & Ross, 2004), and the organization of routines (Feldman & Pentland, 2003) - influence whether actors are able to actualize an affordance.

Simultaneously, calls to examine affordances at the organizational level (Zammuto et al., 2007) and extend theory toward multi-level analyses (Volkoff & Strong, 2017) remain largely unmet. Empirical studies typically analyse these influences as micro-level, in-use conditions - e.g.

whether an individual user can take up an affordance during a task, such as whether an interface is user-friendly (Volkoff & Strong, 2017). However it gives far less attention to the organizational groundwork required for such use to be possible at all (Volkoff & Strong, 2017). In practice, many affordances depend not only on momentary user interaction but on whether the organization has invested in the shared infrastructure, alignment, and preparation that make a technology ready for use in the first place (Besson & Rowe, 2012; Volkoff & Strong, 2017).

3.2.3 Limitations of Affordance Theory for AI agent Research

Despite its strengths, existing affordance research exhibits several limitations that constrain its applicability to emerging forms of agentic, tool-orchestrating AI. First, while scholars emphasize the necessity of analyzing affordances at the organizational level (Volkoff & Strong, 2017), much research focuses on user-technology interactions, offering limited insight into affordances that arise at organizational or interfunctional levels (e.g., Fayard & Weeks, 2014; Keller et al., 2019). Second, the emphasis on during-use dynamics leaves comparatively little theorisation of proactive conditions which may significantly shape whether technological possibilities can be taken up in the first place (Besson & Rowe, 2012; Volkoff & Strong, 2017). This gap becomes particularly relevant for agentic AI systems, which may execute sequences of actions with limited human intervention; in such cases, the organizational and technical environment established before use plays an outsized role in determining what the system can accomplish once deployed. Third, although recent studies identify affordances associated with analytics, predictive models, and generative AI (Bao et al., 2023; Godé et al., 2024; Li, 2025), existing work predominantly examines generative AI rather than the more autonomous, orchestrating capabilities characteristic of agentic AI. As a result, current insights provide limited understanding of how emerging agentic systems may influence innovation activities beyond isolated content generation or analytical augmentation.

Taken together, affordance–actualization theory provides a useful but incomplete foundation for examining emerging forms of agentic AI within innovation processes. Existing research primarily theorises affordances at the level of individual actors and during-use episodes, offering limited guidance on how organizational and technical preconditions shape whether affordances can be realised in the first place. Moreover, prior work has largely focused on predictive or generative systems, leaving underexplored how more autonomous, orchestrating agents interact with multi-stage innovation workflows. This gap is particularly salient given that innovation work involves distributed responsibilities, iterative evaluation, and interdependencies

across functions - conditions under which AI agents may create value only when certain organizational and technical enablers are present.

3.3 Research Gap

Across the four stages of the innovation funnel, existing research provides valuable but fragmented insights into how AI augments discrete innovation activities- from data-driven sensing in Discover, to predictive evaluation in Define, generative design in Develop, and forecasting in Deliver. However, these studies largely examine isolated tools and task-level support, offering limited understanding of how AI reshapes interdependencies, feedback loops, and coordination across the innovation process. The literature overwhelmingly focuses on assistive forms of AI-analytics, prediction, and generative content creation- while systematic evidence on autonomous, multi-step, or orchestrating AI agents is largely absent. Moreover, organizational conditions such as governance routines, data interoperability, workflow design, and cross-functional coordination repeatedly moderate outcomes, yet remain conceptually underdeveloped.

A parallel set of limitations is evident in affordance–actualization research. Prior studies predominantly analyse affordances at the level of individual users and during-use episodes, providing limited insight into the organizational and technical preconditions that shape whether technological potentials can be actualized. Existing work focuses largely on predictive and generative systems rather than agentic AI, leaving unexplored how autonomous, tool-orchestrating systems interact with multi-stage innovation workflows. Building on these insights and answering the call for multi-level inquiry (Volkoff & Strong, 2017; Zammuto et al., 2007), the present study extends affordance–actualization theory by identifying the affordances and value-creation opportunities specific to AI agents and the boundary conditions that shape their realisation across the innovation funnel.

4. Methodology

This chapter outlines the methodological approach used in the study. It first presents the research design (4.1), followed by the procedures for case selection and data collection (4.2). It then describes the coding and analytical strategy used to develop the findings (4.3). Finally, the chapter addresses issues of research quality (4.4) and the ethical considerations that guided the study (4.5).

4.1 Research Design

The research followed a multiple case study design (Eisenhardt, 2021) grounded in constant comparison of theory and data (Glaser & Straulss, 1967). Studying the emerging phenomena of AI agents with the purpose of understanding their potential in innovation processes, compared to confirming it, a qualitative, exploratory and abductive approach was preferred (Yin, 2003). Consequently, we prioritise rich accounts of how organizational members make sense of and enact AI agents in their innovation work over measurement of predefined variables. Given our focus on how organizational members interpret and enact AI agents in practice, the study adopts a constructivist ontology and interpretivist epistemology, treating innovation realities as socially constructed and accessible through actors' meanings rather than objective measurement. The study reviewed Nordic companies of three categories: industry-native, digital-native and AI-native firms, with each category corresponding to a distinct technological basis for their core business and offerings: physical/industrial (non-digital), digital, and AI agent–based technologies.

Table 2 - Case Categories Summary

Category	**No. of Cases**	**Industry**	**Headcount / firm**	**No. of Interviews**
Industry-native	2	Manufacturing	~150-500	8
Digital-native	2	E-commerce	~50-100	10
AI-native	3	AI Software	~3-5	8

To strengthen theoretical rigor in case selection, the three categories were selected using maximum theoretical variation. Industry-native, digital-native, and AI-native organizations represent distinct socio-technical logics with meaningful differences in data infrastructure,

workflow integration, and innovation routines. This variation enables the study to observe both stable patterns across heterogeneous contexts and divergences associated with differing technological and organizational maturities.

4.2 Data Collection

4.2.1 Case Selection

Case selection followed a purposive, theoretically informed sampling strategy (Bell et al., 2022). After initial exploratory interviews, we selected cases that maximised variation across organizational maturity and socio-technical conditions, ensuring that the sample reflected distinct configurations relevant for studying the emergence and realisation of AI agent affordances. Industry-native, digital-native, and AI-native firms thereby provided analytically meaningful contrasts in terms of data architecture, system interoperability, and innovation-process structure. This theoretical heterogeneity formed the foundation for identifying both cross-contextual regularities and context-specific boundary conditions shaping affordance actualization.

One of the unique advantages of including cases of AI-native startups was that we could explore the potential of AI agents in the innovation funnel both from the companies and their customer's point of view - as such adding elements of cross-validation. While the AI-native organizations also provided insights into their customers' innovation conditions, these were not treated as separate cases. Instead, they were used to increase breadth in identifying actualization barriers in contexts where agentic systems are deployed.

Within each category, cases were selected using literal replication logic, meaning that we expected a degree of similarity in patterns recurring among organizations with comparable technological maturity and process structures (Yin, 2003). This logic allowed inclusion of organizations expected to display contrasting organizational and technical boundary conditions relevant for understanding how AI agents interact with different innovation environments. This combination of within-category similarity and across-category contrast created a structured basis for systematic cross-case comparison and strengthened the study's capacity to theorize how organizational context shapes the emergence and realization of value from AI agents.

Cross-case comparability was ensured by selecting cases that all conducted structured innovation activities spanning the innovation funnel stages. Although these activities differed in tempo and tooling across contexts, their functional equivalence enabled structured comparison. The combination of maximum theoretical variation across categories and functional similarity

within cases ensured that observed differences could be attributed to meaningful contextual factors rather than incompatible process structures.

A limitation of comparing different technological archetypes is that AI-native firms are typically young and early-stage, while industry-native firms are often larger and more established. To partially mitigate this, one industry-native and digital-native case founded in the 21st century were selected to reduce the maturity gap and allow a more balanced comparison. Nevertheless, maturity differences remain inherent when studying emerging technologies that diffuse unevenly across organizational contexts, and should be considered when interpreting observed contrasts between categories.

4.2.2 Data Sampling

The level of analysis in this study is the organizational level - that is, the participating case firms. Each case represents an organizational context in which the potential of AI agents can be examined in relation to technological and organizational characteristics. The purpose of the data collection was therefore to assess how the value potential of AI agents varies across firms with different digital and AI maturities and to capture the technical and organizational conditions influencing this potential.

Within each case company, participants were selected to ensure a broad yet analytically coherent coverage of perspectives. Interviews were conducted with senior managers representing seven primary functional domains: C-suite and top management, Technology/IT, Product, Merchandising, Distribution/Operations, Sales and Business Development. These functional groups collectively satisfied the study's selection criteria.

Interviewees were selected using purposive sampling (Patton, 2015) based on four criteria: (1) substantial responsibility for or experience with core operational processes, (2) involvement in innovation or product development processes, (3) responsibility for IT architecture, data infrastructure, or AI initiatives, and/or (4) a senior role providing a holistic view of the organization's innovation activities. This ensured that respondents could provide rich contextual knowledge about both the firm's innovation processes and the technical and organizational conditions that influence the actualization of new technologies.

By combining executive-level respondents with functional experts, the sample captured both top-down strategic perspectives and bottom-up implementation realities. This balance was essential for understanding how AI agents can enhance innovation processes across different organizational layers and decision contexts. In total, 26 semi-structured (Appendix B) interviews were conducted, each lasting approximately 45 to 60 minutes. This composition enabled a

multi-perspective understanding of both value potential and boundary conditions of AI agents across the innovation funnel.

4.2.3 Interview Procedure

Participants were contacted via phone and email through established company channels. Interviews were conducted online to accommodate geographical dispersion and scheduling constraints. All participants received information outlining the study's purpose, confidentiality measures, and data-handling procedures. Anonymity was assured and informed consent was requested for the use of AI-assisted transcription tools.

The interviews followed a semi-structured format guided by a case-study protocol (Appendix A), comprising a general protocol as well as a technical extension for respondents able to provide insight into the organization's IT infrastructure and technical conditions. The general protocol focused on innovation processes, organizational context and perceived opportunities for AI integration. The technical extension gave a majority of its focus to exploring the enterprise and IT architecture. Both protocols were designed based on theoretical and contextual insights from the literature and phrased to minimize bias and encourage open reflection.

Consistent with qualitative case-study methodology, the interview protocol was treated as iterative rather than static. As interviews progressed, emerging themes and patterns were incorporated into subsequent interviews to deepen understanding and refine the direction of inquiry (Yin, 2003). All interviews were recorded and transcribed with transcription software to ensure accuracy and to allow interviewers to remain attentive during conversations. Transcripts were subsequently reviewed and corrected for any transcription errors by the researchers to ensure precision prior to analysis. This systematic and ethically grounded approach ensured that the collected data were both credible and sufficiently rich to capture the multifaceted organizational and technical conditions influencing the actualization of AI agent affordances within each firm.

4.3 Data Coding & Analysis

Our data analysis followed an iterative qualitative strategy combining inductive coding, cross-case comparison, and abductive theorizing. All transcripts and coding materials were organized in a shared workspace, supported by analytic memos documenting emerging insights, category refinements, and coding decisions. We acknowledge that qualitative interpretation is unavoidably shaped by researchers' prior assumptions and theoretical sensitivity (Braun & Clarke,

2006). To mitigate this, we kept an audit trail of evolving coding decisions and revisited early interpretations as new data emerged, helping prevent premature thematic closure.

We began with open coding to surface first-order concepts grounded in participants' language (Gioia, Corley & Hamilton, 2013). These captured expressions of needs, constraints, affordances, and contextual conditions related to the use and perceived value of AI agents. Through repeated reading and recoding, conceptually similar codes were grouped into axial categories (Strauss & Corbin, 1998), which were then consolidated into aggregate dimensions reflecting recurring structures in how organizations understood and enacted AI agents. Saturation was reached when additional interviews no longer generated new first-order codes or altered the category structure.

Once within-case coding stabilised, we conducted constant comparison across cases (Glaser & Strauss, 1967), systematically examining similarities and differences across industry-native, digital-native, and AI-native organizations and across innovation-funnel stages. A cross-case comparison template supported analytical consistency.

Although our coding process was inductive, the subsequent cross-case synthesis drew on a light scaffolding informed by established innovation-process literature and affordance–actualization theory. This structure was applied only after emergent codes and categories had stabilized, enabling us to organize recurring patterns according to (1) stages of the innovation process, (2) the organizational needs and affordances surfaced in each stage, and (3) the boundary conditions that shaped whether this potential could materialize. This approach preserved the grounded nature of the findings while supporting analytical comparability across organizational contexts.

Finally, we engaged in abductive iteration between data and theory (Timmermans & Tavory, 2012). Emerging empirical patterns were continuously confronted with research on affordances and innovation processes, enabling refinement of category boundaries, consolidation of overlapping themes, and synthesis of the aggregate dimensions that underpin the analytical constructs developed. During the later stages of analysis, recurrent interactions between stage-level patterns that extended beyond the boundaries of the funnel appeared. As recommended in abductive research, these unexpected linkages prompted an additional layer of theorising aimed at integrating the emerging insights. This step led to the construction of a higher-order synthesis and model development which captures cross-stage dynamics that were not visible through stage-by-stage analysis alone.

4.4 Quality Considerations

Because qualitative, abductive case research relies on interpretation and contextual understanding, this study assesses rigor using Lincoln and Guba's (1985) trustworthiness criteria. Credibility was supported through triangulation across firm types and roles, iterative refinement of the interview protocol, and manual verification of AI-generated transcripts. Transferability was enabled by maximum theoretical variation in case selection and rich, contextually grounded descriptions of each firm's socio-technical conditions and innovation practices, in line with Geertz's (1973) notion of "thick description". Dependability was strengthened through transparent documentation of coding decisions, analytic memoing, and a structured cross-case comparison process. Confirmability was strengthened by grounding higher-level interpretations in first-order data and revisiting category boundaries during abductive iterations. To reduce individual researcher bias, we divided initial coding responsibilities across cases, independently coded overlapping excerpts, and reconciled discrepancies through discussion, supported by regular supervisory feedback. Theoretical scaffolding was applied only after data-driven patterns had stabilised. Together, these measures provide a traceable and analytically cautious foundation for the study's insights into how agentic AI interacts with organizational innovation processes.

4.5 Ethical Considerations

Participants received clear information about the study and gave informed consent before interviews were recorded. Organizational and individual identities were anonymised to protect confidentiality. Recordings were deleted after transcription, and AI-assisted transcripts were anonymized and manually reviewed to ensure accuracy. These measures safeguarded participant privacy and data integrity throughout the research process.

5. Findings

This section presents the empirical findings of the study. It begins by outlining the three organizational contexts and how innovation work unfolds within each case category (5.1). It then compares recurring and divergent patterns across these contexts (5.2). Finally, it examines how AI agent affordances manifest across the stages of the innovation funnel, analysing the conditions that enable or constrain their actualization and how these vary across organizational environments (5.3).

5.1 Case Overviews

5.1.1 Industry-native

The interviews from the industry-native organization revealed six recurring patterns that describe how innovation work unfolds in this context. The first concerned fragmented data architecture and missing semantic integration. Interviewees consistently described difficulties accessing and consolidating information, noting that data were scattered across legacy systems, spreadsheets, and individual expertise. This fragmentation required employees to spend substantial time locating, verifying, and reconstructing information and context before they could proceed with exploration or evaluation tasks.

"We have some problems with the registers not always being harmonized… the same attribute can be called 1072 in one system and something entirely different in another. It becomes difficult to match without building a translation register [...] we often end up piecing things together from old files, emails, and scattered systems because nothing is linked all the way through."

- Market Developer, Industry-native #1

A second pattern reflected resource-intensive innovation processes with overloaded evaluation capacity. Participants explained that activities such as pre-studies, feasibility assessments, and responses to Requests for Quotation (RFQ) were highly manual and time-consuming. While ideas and opportunities were generated frequently, the capacity to assess them was limited. Several interviewees commented on the backlog of potential initiatives that remained unevaluated due to time constraints.

A third theme involved the scattered nature of innovation discovery. Innovation was described as arising from a variety of sources and business development functions often served as analytical hubs, interpreting inputs and coordinating early-stage activities.

"New business opportunities can be driven by regulatory requirements, shortages of raw materials, or that we want to start producing a product ourselves. We get ideas that are customer-driven or from ourselves […] we

look at import statistics, who produces the product today, which customers buy it, the volumes, and whether it fits our terminal strategy."

- Market Developer, Industry-native #1

A fourth theme was the high perceived potential for agentic AI across multiple parts of the innovation process. Interviewees expressed strong expectations about how AI agents could support problem discovery and analysis of different solutions. Many participants noted that substantial efficiency gains were possible if information were better harmonized and more easily accessible. Interviewees highlighted the potential of multi-agent workflows, greater speed in information retrieval, and the ability to automate early-stage scanning.

A fifth recurring pattern was the presence of a supportive organizational culture but insufficient structural readiness for agent adoption. Interviewees described the culture as open, curious, and positive toward exploring new digital technologies. At the same time, they emphasized structural limitations such as unclear governance and ownership of digital initiatives, as well as limited internal resources. These gaps created uncertainty about how agent-based solutions would be integrated or maintained in practice.

"We built [the company] from nothing with limited resources … our priority was to build a disruptive technology and lead a fundamental shift in the sensor technology industry, not creating the best business operating system. [Therefore], our organization was designed to drive innovation within our hardware and software-on-chip solutions - governance and digitalization was secondary. Now, as a mature publicly listed company, we have shifted more attention to address this digitalization debt and see a lot of potential in AI across our departments."

- CEO, Industry-native #2

Finally, the interviews highlighted safety, compliance, and domain-specific constraints requiring human judgment. Participants noted that while early-stage scanning, exploration, and evaluation could be meaningfully supported by automation, decisions involving functional safety, regulatory compliance, or customer-specific formulations required expert interpretation. Automation was therefore seen as complementing, rather than replacing, critical human oversight.

5.1.2 Digital-native

The interviews from the digital-native organization revealed six recurring patterns that characterize how innovation work is carried out in this context.

A prominent theme concerned data infrastructure constraints. Interviewees described a landscape in which information relevant for innovation was present in numerous systems but was difficult to access, aggregate, or interpret. Several participants noted that data were often

underused, that insight generation was highly manual, and that forecasting or merchant-related understanding required extensive rework. Low data literacy in parts of the organization further complicated the process of drawing reliable conclusions from available information.

"Frontend [E-commerce] operations rely on scattered app integrations and data sources to build [state-of-the-art] features that optimize [sales] conversion […] we lack the ability to handle the complexity of our own [technical] setup or make use of our own analytics data [...] so we rely heavily on external agencies to improve the customer experience."

- Chief Technology Officer, Digital-native #2

A second theme related to innovation process immaturity and structural weaknesses. Participants described prioritization as slow and bottlenecked, with evaluation processes that were largely ad hoc rather than systematic. Decision-making was frequently top-down, and employees noted the absence of a structured innovation funnel or consistent mechanisms to support bottom-up idea development. This resulted in uneven progression of innovation tasks and uncertainty about how initiatives should move forward.

A third theme involved a pattern of operational firefighting systematically limiting organizational capacity for innovation efforts. Participants frequently noted that teams were forced to address immediate operational issues at the expense of strategic development, making it difficult to allocate the sustained attention and cross-functional time necessary for exploratory work. These resource limitations and cross-functional dependencies meant that innovation perpetually lost out when operational pressures intensified.

"We have more ideas than tech capacity […] so it's always hard prioritization. We've had a peak of fires because of recent platform migration - everything else has had to wait until after Black Week."

- Chief Technology Officer, Digital-native #1

A fourth pattern concerned fragmented technological environments and missing orchestration. While digital tools were abundant, they were distributed across different teams with limited interoperability. Interviewees described constraints in technical experimentation, gaps in governance and safety practices, and the absence of any orchestration layer that could coordinate or integrate tools across departments. This fragmentation created friction when attempting to test or scale new technical solutions.

A fifth theme reflected emerging AI adoption and capability building. Interviewees reported rising interest and literacy in AI, with early experiments occurring in customer service, operations, and creativity-related tasks. Some automations were already in place, though typically unscaled or isolated. Participants also described growing curiosity about agents for daily routines or task support, but noted that usage varied widely across the organization

Finally, the interviews revealed a theme related to leadership signals and cultural dynamics. Interviewees described strong enthusiasm for innovation from senior leaders but noted that this sometimes translated into top-down cultural expectations rather than broadly distributed empowerment. In addition, learning diffusion was often reported as uneven.

5.1.3 AI-native

The interviews across the AI-native organizations revealed seven recurring patterns that describe how innovation processes work in teams that build agentic AI systems internally and deploy them to customers.

A prominent theme concerned agentic automation fundamentally restructuring internal product development velocity. Interviewees consistently emphasized that agents reshaped internal workflows by collapsing development, design, testing, documentation, and go-to-market preparation into far shorter cycles. Several participants described how tasks that previously moved sequentially across functions could now be completed in parallel or near-simultaneously, producing a qualitatively different pace of innovation. Teams repeatedly noted that agent-driven workflows compressed timelines by automating multi-step tasks that historically involved several handovers.

"[Our AI agent operating system] is dynamic, multimodal, and responsive in real-world environments. Guests see instant results on-screen maps, menus, bookings, or stay details, while the agent can also call a restaurant, notify a shuttle driver, or trigger a maintenance task."

- Co-founder, AI-native #2

A second recurring pattern reflected how customer's realization of agentic innovation depends on guided onboarding, education, and cultural readiness. Interviewees noted that successful adoption required more than just access to tools: it depended on structured onboarding sessions, early-stage handholding, shared experimentation norms, and an explicit culture that supported learning and curiosity. Customers that adopted agents most successfully had routines for collaborative experimentation and reflection.

A third theme involved how technical integration foundations determine the customer's ceiling of agentic affordances. Interviewees described the technical landscape where systems and software tools formed the infrastructural backbone needed to unlock more advanced types of automation. Participants also emphasized that integration is valuable primarily because agents depend on reliable, richly connected context to reason effectively.

"Across [agentic AI] use cases, it usually comes down to context - finding the right context and enabling in-context reasoning. This means connecting to [the customer's] preferred apps and systems and extracting the

relevant information. The hardest part for the agent isn't executing the task, it's understanding why it's being done and how it fits into everything else."

- Technical co-founder, AI-native #3

A fourth theme centered on the staged maturity path of agentic AI realization in innovation processes. AI-native organizations consistently observed that customers began with relatively simple use cases such as search or consolidation, then progressed toward coordination support, and only later reached automation that spanned multiple systems or departments. This progression was described as gradual, with each stage dependent on the customer's underlying process and system maturity. Interviewees highlighted that these preconditions are what allow customers to move from simple retrieval to higher-order automation.

A fifth pattern described how agentic systems surface hidden innovation opportunities by unifying customer's internal and external data. Interviewees explained that agents were able to integrate communication logs, customer feedback, sales information, and relevant external sources in ways that exposed patterns not previously visible. Several participants highlighted examples where agents uncovered surprising insights by combining data that had historically lived in different systems or functions. The ability to unify distributed information resonated with broader accounts of agents helping users navigate complexity across fragmented digital environments.

"The core value of our product comes from helping people manage complexity, especially when digital information, communication, and decisions are spread across many channels and the information load makes it hard to navigate between what's important and what's simply noice"

- Co-founder & CEO, AI-native #3

A sixth theme concerned how market readiness and organizational maturity limits AI agent adoption for customers. Interviewees emphasized that the extent to which customers could adopt agentic solutions depended heavily on factors such as their tool ecosystem readiness, degree of digital maturity, willingness to experiment, and the prevailing mindset toward augmenting existing workflows. Differences in these readiness factors often produced wide variation in how quickly and whether or not customers can benefit from AI agent affordances.

Finally, the interviews highlighted how agents act as cross-team cognitive integrators by aligning divergent perceptions. Interviewees noted that agents frequently synthesized multi-perspective signals revealing underlying misalignments in reasoning or priorities. This integration enabled more coherent and aligned decision-making across teams, especially in complex or rapidly evolving projects.

5.2 Cross-Case Empirical Findings

Across the three organizational contexts, several empirical similarities emerged regarding how innovation work unfolds in practice. First, interviewees in all cases described challenges related to accessing, consolidating, and making sense of information needed for exploration, evaluation, or decision-making. Despite differences in their technological landscapes, all organizations experienced some form of data fragmentation that required employees to manually reconstruct context before progress could be made. This often involved searching across multiple systems, validating details, or relying on colleagues with specific knowledge.

A second shared pattern concerned the coordination demands involved in progressing innovation tasks. Participants across cases spoke about repeated check-ins, clarification loops, or alignment activities necessary to ensure mutual understanding among different roles or teams. Although the underlying reasons varied, the effect was consistent: coordination absorbed time and delayed the flow of innovation work.

Interviewees in all organizations also noted that innovation efforts were frequently interrupted by operational demands, shifting priorities, or external pressures. This created stop-start dynamics that made it difficult to maintain momentum across longer innovation cycles. Participants described how exploratory tasks were often paused or deprioritized in favour of more urgent day-to-day responsibilities.

Finally, participants across all contexts expressed interest in more advanced digital or AI-based support tools, even though current use remained limited. Interviewees consistently conveyed curiosity and optimism about the potential benefits of such tools, while simultaneously acknowledging that adoption was still in early stages.

Despite these similarities, clear differences emerged in the sources and expressions of innovation-related friction. In the industry-native context, interviewees described a more cautious but pragmatic openness toward AI support while emphasizing that adoption is inherently constrained by the domain in which they operate. The core barriers were tied to legacy systems, strict regulatory requirements, and a heavy reliance on tacit, experience-based judgment. Safety and compliance considerations defined not only what could be automated, but also the pace at which experimentation was perceived as feasible.

In the digital-native context, friction was driven more by the pace and volume of ongoing operations. Interviewees described a fast-moving environment with shifting priorities, fragmented tooling across functions, and limited structure for bottom-up idea development. Although more digitally mature than the industry-native cases, this environment was

characterised by ad hoc evaluation routines, cross-functional dependencies, and difficulty securing focused time for exploration.

The AI-native organizations differed from both other categories in that internal innovation workflows were described as stable, integrated, and supported by consistent routines for experimentation. Interviewees portrayed their internal processes as relatively smooth compared to those of their customers, whose system maturity, data environments, and readiness for experimentation varied considerably. As a result, friction for these organizations arose primarily when deploying solutions externally rather than internally.

5.3 Cross-Case Comparison - Affordances across the Innovation Funnel

Building on the individual case overviews and category summaries in Sections 5.1 and 5.2, this section integrates the findings into a cross-case comparison. This aims to present how AI agent affordances manifest across the four stages of the innovation funnel and how these affordances interact with stage-specific organizational needs.

The subsequent tables (3 through 7) map the distinct innovation requirements of each case category to the agentic affordances relevant at that stage. They also indicate the observed maturity of these affordances within each organization, distinguishing between absent, emerging and present forms. Finally, the tables outline the boundary conditions that shape and constrain actualization, highlighting the contextual factors that either enable or limit the use of agentic AI across stages.

5.3.1 Discover

Table 3 - Discover

Case Category	**Organizational Need**	**Affordance Identified**	**Presence Level**	**Boundary Conditions**
Industry-native	Unified access to scattered internal knowledge	Copilot-assisted document search (Quality, safety & product info)	Emerging	Fragmented systems; inconsistent IDs; local-only search; no semantic layer

	External market, regulatory, supplier opportunity sensing	Agents for regulatory scanning, external news and database scanning	Identified	No external-data integration; manual Excel extraction; no unified access point
	Early multi-source feasibility insight	Agentic feasibility snapshot (CRM + supplier + import data unified)	Identified	No semantic integration; no cross-system connectors; manual data assembly
Digital-native	Understanding traffic shifts, merchant behavior, external signals	Agent-assisted external signal interpretation (search trends, competitor monitoring)	Emerging	Low analytical capacity; fragmented tools; inconsistent signals; limited merged data layer
	Faster insight generation from internal marketplace data	Agent-assisted insight synthesis (categorization logs, delivery patterns, pricing signals)	Identified	Data fragmentation; low data literacy; tool immaturity; siloed systems
	Competitor and market pattern monitoring	Agent-based scraping (used by small pockets)	Emerging	Not integrated; no orchestration; unclear ownership
AI-native	Faster extraction of signals from internal product and communication data	Internal communication tools scraped for patterns and themes	Present	Requires high-quality indexing (fulfilled internally)
	Discovery through automated UI/code exploration	Agents run code exploration + documentation parsing	Present	Requires integrations (fulfilled)

Across cases, discovery work is universally characterised by high information demands, but the maturity and actualization of affordances vary sharply. In the industry-native organization, discovery depends on heterogeneous inputs spread across legacy systems, resulting in limited actualization of identified affordances such as document search, regulatory scanning,

and feasibility snapshots. Fragmented systems, inconsistent identifiers, and missing semantic layers prevent these affordances from materialising in practice.

The digital-native organization operates on fast-moving commercial signals but struggles to synthesize them due to fragmented analytics and limited ability to filter noise. Affordances related to external-signal interpretation and insight synthesis are identifiable but remain unrealised because of weak interoperability, uneven analytical capacity, and distributed toolchains.

In contrast, the AI-native organizations exhibit significantly more mature discovery practices. Since agentic AI features are core to their product offering, AI-native firms can also deploy them internally, enabling affordances such as knowledge retrieval, communication-pattern extraction and exploratory data analysis. While internal affordances are largely present, external affordance remains constrained by customers' lower technical maturity and fragmented data environments.

Overall, the Discover stage reveals a consistent pattern: while all organizations express strong discovery needs, realized agentic affordances emerge only where data integration, indexing, and interoperability are sufficiently developed. This divergence frames the comparative analysis for the subsequent stages of the innovation funnel.

5.3.2 Define

Table 4 - Define

Case Category	Organizational Need	Affordance Identified	Presence Level	Boundary Conditions
Industry-native	Structured evaluation of ideas & pre-studies	Agents synthesizing and importing data; regulatory information in evaluation briefs	Identified	Semantic integration missing; data harmonization absent; experimentation with agents
	Interpret feasibility inputs faster	Agents clustering & summarizing multi-source technical data	Identified	Manual data prep needed first; no unified dashboards

Digital-native	Evidence-based prioritisation & idea evaluation	Agents generating draft specs, feasibility checks, capacity-aware estimates	Identified	Low data literacy; inconsistent routines; no unified evaluation system
	Faster qualitative/quantitative insight synthesis	AI-assisted synthesis of user research, merchant signals	Identified	Fragmented inputs; no merged data layer; bottlenecked analysts
AI-native	Rapid evaluation of internal product hypotheses	Agent-supported feasibility checks, UI iterations, documentation drafts	Present	Requires prompt libraries & internal experimentation culture
	Cross-functional alignment around problem framing	Agents summarising differing signals & aligning mental models	Present	Requires coordination maturity (fulfilled internally)
	Early customer idea evaluation	Clustering customer feedback, summarizing demand themes	Emerging	Customer data hygiene & semantic integration vary

Across cases, the Define stage centres on making sense of diverse inputs and evaluating early ideas, yet actualization of affordances diverges sharply. In the industry-native organization, evaluation relies on assembling fragmented technical, regulatory, and customer information. Identified affordances focus on consolidating feasibility inputs and generating structured evaluation briefs, but data fragmentation, missing semantic integration, and manual preparation requirements prevent meaningful adoption.

The digital-native organization faces a high volume of ideas and strong commercial pressures, but evaluation routines remain inconsistent and constrained by distributed tools and uneven data literacy. Affordances for generating draft specifications, synthesising qualitative and quantitative signals, and supporting prioritisation are identifiable but cannot be realised due to fragmented inputs, limited analytical capacity, and the absence of an integrated evaluation system.

The AI-native organizations exhibit the most mature Define-stage affordances. Agents already assist with feasibility checks, documentation drafts, UI iterations, and cross-functional alignment. Several affordances are present internally, supported by a culture of experimentation and strong coordination routines. Externally, similar affordances are emerging for customers, though actualization depends on client-side data hygiene, connector availability, and organizational readiness.

Overall, the Define stage shows broad convergence in needs - structured evaluation, insight synthesis, and prioritisation - but persistent divergence in actualization. Industry- and digital-native firms are limited by fragmentation, capacity constraints, and weak integration, while AI-native firms demonstrate that agentic support becomes viable once foundational conditions are met. These contrasts frame the next stage of the comparative analysis.

5.3.3 Develop

Table 5 - Develop

Case Category	**Organizational Need**	**Affordance Identified**	**Presence Level**	**Boundary Conditions**
Industry-native	Automating pre-study technical work	Agents preparing requests for quotation, supplier lookups and cost models	Identified	System interoperability missing; master data not harmonized
	Chemical quality anomaly detections	Agent-enabled anomaly spotting in lab environments	Identified	Data not unified; requires human (high risk of chemicals)
Digital-native	Accelerating developer & designer workflows	Copilot-type assistance in cod/testing/ prototyping	Present	Fragmented tool ecosystem; limited sandbox; uneven adoption
	Workflow automation within certain ops teams	Existing agents for categorization, grouping, compliance	Emerging	Cannot scale due to lack of orchestration; unclear ownership

AI-native	End-to-end acceleration	Automated UI generation, testing, code writing	Present	Requires orchestration and strong connectors
	Cross-role augmentation	Agents assisting across all functions	Present	Requires guided onboarding and cultural norms

In the industry-native organization, the Develop stage remains dominated by resource-intensive pre-study work and technical feasibility assessment. Identified affordances focus on automating these discrete tasks and supporting anomaly detection, yet none are realized due to fragmented systems, missing semantic links, lack of APIs, and domain constraints that require extensive human judgement.

The digital-native organization shows broader experimentation, especially in engineering and operations. Agents support code generation, review, testing, and selective operational automations, but adoption is uneven and localized. Fragmented tool environments, limited orchestration, and unclear ownership prevent these affordances from scaling into a coherent development workflow.

The AI-native organizations demonstrate the most advanced Develop-stage affordances. Agents automate large components of the development pipeline, including UI generation, testing, documentation creation, and iterative design or engineering support. Strong integration across tools enables cross-role augmentation and efficient orchestration. Externally, similar affordances are emerging for customers but remain constrained by client-side data quality, connector availability, and governance readiness.

Overall, the Develop stage shows the widest gap in agentic maturity across cases. All organizations express similar needs for accelerated iteration, reduced manual overhead, and better integration of technical artefacts, but only the AI-native organizations possess the conditions required for actualization of internal affordances. However, AI-native organizations also noted that customer readiness can constrain how far these affordances translate into their external product offering.

5.3.4 Deliver

Table 6 - Deliver

Case Category	Organizational Need	Affordance Identified	Presence Level	Boundary Conditions
Industry-native	Scalable implementation and deployment of validated innovation outputs	None identified	Absent	Upstream constraints dominate; no downstream capacity
Digital-native	Operationalize or scale validated innovation outputs across teams	None identified	Absent	Bottlenecks in data integration + governance prevent downstream progress
AI-native	Seamless flow and automated handoffs from development artifacts to delivery/ documentation.	Automated documentation synthesis and deployment transition support.	Emerging	Requires unified metadata and version control systems

In industry-native organizations, Deliver activities remain tightly bound to regulatory validation, safety requirements, and cross-departmental coordination. These compliance-heavy workflows limit the scope for automation, and fragmented systems prevent agents from supporting downstream tasks even where potential exists.

In digital-native settings, teams use isolated automations within operations, but these are disconnected from the deployment of innovation outputs. Structural issues such as weak orchestration, unclear ownership, and siloed tooling prevent agentic workflows from extending into coordinated implementation processes.

The AI-native organizations show the closest approximation of Deliver-stage augmentation. Strong internal integration has allowed agents to start bridging the Develop-Deliver boundary by generating technical documentation and preserving context from upstream work. However, these affordances remain general workflow support rather than fully developed Deliver-specific capabilities, and external actualization is limited by customer-side constraints such as inconsistent metadata and version control practices.

Overall, the Deliver stage reveals a consistent structural asymmetry: while upstream stages exhibit clear agentic potential, downstream implementation remains constrained by

regulatory, organizational, and integration barriers. Even in the AI-native case, agentic support is limited to context transfer rather than autonomous deployment or scaling, highlighting Deliver as the least transformed frontier of the innovation process.

5.3.5 Cross-stage patterns

Table 7 - Cross-stage patterns

Case Category	Organizational Need	Affordance Identified	Presence Level	Boundary Conditions
Industry-native	Reduce context loss across functional silos and document chaos.	Natural-language query across all internal knowledge sources.	Identified	Missing system connectors
	Ensure structured handoffs from discovery to development.	Context-preserving funnel memory.	Identified	IT capacity limitations; missing interoperability/ orchestration framework.
Digital-native	Coordinate tasks and insights across fragmented product and operations teams.	Agentic coordination layer unifying workflows	Identified	No orchestration layer; unclear ownership and governance; siloed systems.
	Reduce manual overhead across repetitive cross-functional routines (e.g., compliance, grouping).	Scaled local automations extending to cross-functional workflows.	Emerging	Fragmented systems prevent execution; automations exist but are unscaled.
AI-native	Accelerate end-to-end innovation cycles by removing handoffs.	Autonomous Code-to-Deployment transitions	Present	Needs unified metadata & version control (internal constraint).

Enable customers to scale from simple search to complex coordination.	Staged innovation enablement	Emerging	Customer maturity varies; integration gaps/lack of customer governance.

In the industry-native organization, cross-stage affordances remain purely aspirational. Participants articulated a clear need for context-preserving mechanisms, such as multi-agent chains linking trend detection, evaluation, and documentation. However, severe fragmentation across systems, combined with limited integration capacity and no semantic layer, prevents even basic cross-stage handoffs from being automated or augmented. Across cases, the lack of stable routines for making sense of and consolidating prior work also prevented upstream rationale from systematically carrying forward between stages.

The digital-native organization shows stronger awareness of the need for cross-stage coordination, driven by operational tempo and dispersed ownership. Participants pointed to the potential of an agentic coordination layer to unify workflows spanning tools such as ticketing, repositories, and communication systems. Yet the absence of orchestration infrastructure, ownership clarity, and integrated governance means that existing local automations cannot scale into sustained, stage-spanning workflows.

The AI-native organizations exhibit the most advanced cross-stage affordances within their internal environment. Agents already automate transitions between engineering, product, and documentation workflows, enabling continuous context retention across develop-to-deliver activities. This demonstrates the highest observed ceiling for cross-stage agentic actualization. However, extending these affordances to customers remains constrained by client-side gaps in data maturity, interoperability, and governance structures.

Overall, the cross-stage findings indicate that while the need for continuous knowledge transfer is universal, the ability to realize high-order agentic continuity depends entirely on the alignment of organizational routines, metadata structures, and integration foundations established upstream. Cross-stage value is therefore not a function of organizational category but of cumulative readiness across the funnel.

6. Analysis

This section develops the analytical interpretation of the study by moving from empirical patterns to a conceptual explanation of how organizations perceive and realise value from AI agents. The analysis proceeds in three steps. First, we introduce the emergent theoretical constructs derived from the data (6.1), comprising the affordance potentials that describe the forms of value agentic AI can enable and the actualization conditions that shape whether this value materialises. Second, we examine how these constructs interact across the stages of the innovation funnel (6.2), clarifying why similar technological capabilities yield different outcomes across contexts. Finally, we integrate these insights into the Agentic Value Ladder (6.3), a mid-range model that explains how organizations progress from lower to higher order agentic value creation as their technical and organizational foundations mature. Taken together, these steps provide the basis for theorising how AI agents reshape innovation processes across different organizational environments.

6.1 Emergent Theoretical Constructs

6.1.1 Affordance Potentials

The concept of Affordance Potential is introduced here as an emergent theoretical construct to capture the precise intersection where technological capacity converts into organizational value. While standard affordance theory defines an affordance as the alignment of a user's action goal with a technical feature, this study argues that an action possibility alone does not guarantee utility. An Affordance Potential, therefore, is conceptualized as the combination of a real organizational need and an enabling affordance that allows that need to be met. By anchoring the affordance in a validated business requirement, this construct presents not merely an abstract action possibility, but a theoretically grounded value creation opportunity. Across the cases analyzed, the synthesis of recurring organizational needs with their corresponding matching affordances reveals the five distinct Affordance Potentials detailed below. These findings indicate a spectrum of value creation: while the first potential addresses the acceleration of existing workflows (efficiency), the subsequent four unlock qualitatively new capabilities for coherence, exploration, and coordination (novelty).

AP1 - Local Agentic Acceleration denotes the capacity of agents to autonomously execute knowledge-intensive tasks - such as pre-studies, feasibility checks, coding, testing, and documentation - with minimal human intervention. It addresses a critical organizational need to reduce manual workload and eliminate the repetitive reconstruction of context. Crucially, unlike

subsequent potentials that focus on generative divergence or strategic novelty, Agentic Acceleration is grounded in operational efficiency: it creates value by increasing the velocity and autonomy of existing workflows rather than by introducing novel types of output.

AP2 - Data-to-Insight Leverage captures the potential for agents to convert dispersed internal and external information into structured insights, patterns, and explanations. It reflects the need for coherent sensemaking in environments characterized by fragmented data and scattered knowledge.

AP3 - Cognitive Integration and Coordination refers to the potential for agents to maintain shared context across teams, align interpretations, and surface misalignments. This addresses the coordination challenges observed when multiple roles contribute to evaluation, refinement, and development activities.

AP4 - Opportunity Surface Expansion denotes the potential for agents to broaden the scope of opportunity and idea identification and evaluation by integrating signals from external sources, historical data, feedback, and market patterns. In this sense, they function as a higher order application of data-to-insight leverage, enabling a firm to expand the amount of opportunities identified and effectively turned into ideas. This addresses the reactive nature and limitations of effective scanning capacity within discovery work.

AP5 - Cross-Stage Continuity captures the potential for agents to preserve assumptions, rationale, and accumulated knowledge across the innovation funnel, not just within specific stages. It reflects the widespread issue of context loss between discovery, definition, and development, which leads to rework and incoherent progression.

6.1.2 Actualization Conditions

While the affordance potentials describe what value a new technology like AI agents could create, the findings made clear that its realization was far from automatic. Across the material, participants repeatedly pointed to a set of technical (T), organizational (O), and domain-specific (D) circumstances that shaped whether an affordance progressed, stalled, or never emerged. Through iterative comparison across cases, these recurring patterns crystallised into what we term actualization conditions - the empirically observed organizational factors that either enabled or constrained the value potential of agentic systems:

AC-T1 - Data and Semantic Integration refers to the degree to which relevant information is accessible, connected, and represented in a consistent semantic structure. Fragmented data landscapes in the industry- and digital-native organizations prevented agents from generating

integrated insights or maintaining cross-stage context. Even in AI-native settings, limited customer-side data integration remained a primary constraint.

AC-T2 - Orchestration and Interoperability denotes the extent to which tools, systems, and workflows can interoperate and support multi-step or cross-functional processes. Weak orchestration restricted the feasibility of agent-supported coordination and continuity, particularly where tool ecosystems were siloed or lacked shared protocols.

AC-O1 - Agentic Literacy and Experimentation Culture represents the degree to which individuals understand how to work with agentic systems and engage in iterative experimentation. Variation on this dimension influenced whether teams could integrate agents into existing routines.

AC-O2 - Evaluation Process Capacity captures the time, attention, and structured exploratory routines required to recurrently use agents for opportunity identification and evaluation. In organizations with weak or irregular innovation practices, limited slack and absent evaluative structures meant that agents were not systematically deployed for exploratory assessment, constraining the organization's ability to surface, interpret, evaluate and act on opportunities.

AC-O3 - Governance and Ownership Clarity refers to the clarity of responsibilities, decision rights, and oversight for agentic initiatives. Ambiguity around ownership slowed adoption, limited cross-functional integration, and constrained affordances that depend on stable alignment.

AC-D1 - Domain Constraint Intensity captures the regulatory, safety, contractual, and compliance conditions inherent to a given domain that set the structural ceiling on feasible agent autonomy. While this ceiling constrains how far agentic capabilities can be actualized - as seen in chemical and industrial settings where safety rules strictly limit autonomous decision-making - it also operates as a generative boundary. In several cases, the very intensity of domain constraints created acute operational pressure that motivated early agentic innovations. In e-commerce, for example, compliance obligations and data-quality requirements produced overwhelming manual work that directly triggered the pursuit of AP1 (Agentic Acceleration). Similarly, in B2B SaaS and AI-provider contexts, contractual liability and intellectual-property risks drove the development of verification and guardrail systems. AC-D1 thus both delimits the upper bound of agent autonomy and shapes which affordances become prioritized in practice.

Together, these conditions explain why affordance potentials vary in their degree of realization. They clarify why advanced affordances require more demanding combinations of

technical and organizational foundations, and why gaps in these foundations systematically prevent progression across the cases.

Across the three case categories, the interplay between affordance potentials and actualization conditions produces distinct tension profiles. In industry-native organizations, the strongest needs concern knowledge integration and coordination, yet foundational data structures and governance routines are weakest, creating a structural gap between need and feasibility. In digital-native organizations, advanced tooling and analytical capability coexist with fragmented workflows and dispersed ownership, producing a maturity paradox in which digital capability does not translate into affordance realization. In AI-native firms, internal agentic capabilities are comparatively advanced, but value realization depends on customer environments that often lack equivalent technical or organizational readiness. This creates an asymmetry in which internal maturity exceeds external actualization conditions. Together, these tensions explain why organizations diverge in realized agentic value despite facing similar upstream innovation challenges.

6.2 Affordance Potential and Actualization Conditions across the Innovation Funnel

Table 8 presents a AP–AC-Funnel matrix, linking each affordance potential to the technical and organizational conditions required for its realization across the innovation funnel. This matrix provides the analytical scaffold for the section by clarifying which affordance potentials are theoretically available at each stage and which conditions must be in place for them to materialize. The stage-by-stage analysis below follows this structure, detailing (1) the relevant affordance potentials, (2) how they manifested empirically across cases, and (3) the enabling or constraining conditions observed. This creates direct traceability between the findings, the emergent constructs introduced in Section 6.1, and the cross-stage logic later synthesized in the agentic value ladder.

Table 8 - AP–AC-Funnel Matrix

Innovation Stage	**Affordance Potential (AP)**	**Actualization Conditions (AC)**	**Stage Example**
Discover	AP1 - Local Agentic Acceleration	Primarily AC-O1 and AC-D1, minimal impact from AC-T1	Automating opportunity research

	AP2 - Data-to-Insight Leverage	Primarily AC-T1 but enabled even further by AC-T2 (Basic orchestration)	Synthesizing multi-source data into opportunities
	AP4 - Opportunity Surface Expansion	AC-T1, AC-T2, AC-O1, AC-O2	Scanning broadly for non-obvious opportunities
Define	AP1 - Local Agentic Acceleration	AC-O1, AC-D1	Speeding up feasibility assessments
	AP2 - Data-to-Insight Leverage	AC-T1, AC-O1, AC-D1	Combining internal and external signals into feasibility assessments, trade-off analyses, structured summaries
	AP3 - Cognitive Integration & Coordination	AC-O1, AC-T1, AC-T2, AC-D1	Aligning differing assumptions across teams during evaluation; reducing cycles of misalignment
	AP4 - Opportunity Surface Expansion	AC-T1, AC-T2, AC-O1, AC-O2	Coherently translates insights to ideas and evaluates them
Develop	AP1 - Local Agentic Acceleration	AC-T2, AC-O1, AC-D1,	Automating iteration routines (e.g. documentation, approval processes)
	AP3 - Cognitive Integration & Coordination	AC-O1, AC-T1, AC-T2, AC-D1	Maintaining shared context across teams and surfacing inconsistencies in assumptions

Deliver	(No discernible APs identified)	-	No stage-specific affordance potentials emerged empirically
Cross-Stage	AP5 - Cross-Stage Continuity	AC-T1, AC-T2, AC-O1, AC-O2, AC-O3, AC-D1	Maintaining rationale, assumptions, and historical context across stages

Across cases, fragmented information environments and irregular exploratory routines shaped which affordances could activate in the Discover stage. Three affordance potentials appeared relevant: Local Agentic Acceleration (AP1), Data-to-Insight Leverage (AP2), and Opportunity Surface Expansion (AP4). AP1 was universally recognized but rarely realised because organizations lacked both the literacy and experimentation norms (AC-O1) and the minimal semantic integration (limited AC-T1) required for even basic automated structuring of early opportunity inputs. Domain constraints (AC-D1) further limited where exploratory automation was permissible, particularly in regulated industry-native settings. AP2 occurred only in AI-native contexts, where stronger semantic integration (AC-T1) and basic orchestration (AC-T2) allowed agents to combine conversational data, telemetry signals, and qualitative feedback into unified early "opportunity snapshots"; in contrast, industry- and digital-native organizations relied on disconnected spreadsheets, shared drives, and local knowledge, preventing meaningful synthesis. AP4 was aspirational across all firms: participants described needs for broader scanning of competitor, regulatory, and customer signals, but this affordance requires a combination of semantic integration (AC-T1), orchestration (AC-T2), literacy norms (AC-O1), and evaluation capacity (AC-O2), none of which were sufficiently established. As a result, Discover is bounded by low data coherence, limited organizational experimentation, and weak orchestration, leaving most affordance potentials latent rather than realised.

In the Define stage, organizational conditions became especially consequential for affordance realisation. Four affordance potentials were relevant: Local Agentic Acceleration (AP1), Data-to-Insight Leverage (AP2), Cognitive Integration & Coordination (AP3), and Opportunity Surface Expansion (AP4). AP1 was widely desired - for instance, to structure feasibility snapshots or synthesize technical inputs - but its impact was constrained by limited literacy and experimentation norms (AC-O1) and, in several industry-native settings, domain restrictions (AC-D1) that constrained more autonomous evaluation. AP2 again appeared only in AI-native organizations, where semantic integration (AC-T1) and literacy norms (AC-O1), together with favourable domain conditions (AC-D1), enabled agents to combine internal and

external signals into structured feasibility assessments; elsewhere, fragmented and poorly indexed inputs prevented such synthesis. AP3 was the most advanced affordance in this stage but also the least achievable: aligning assumptions across product, engineering, sales, and compliance requires literacy norms (AC-O1), semantic integration (AC-T1), orchestration (AC-T2), and favourable domain contexts (AC-D1), a configuration observed only in AI-native environments. Finally, AP4 - translating insights into alternative problem framings or strategic scopes - depended on semantic integration (AC-T1), orchestration (AC-T2), literacy (AC-O1), and evaluation capacity (AC-O2), but these capabilities rarely coexisted outside AI-native settings. Define therefore illustrates that even when technical maturity improves, organizational conditions remain the primary bottleneck to realising higher-order affordances.

The Develop stage exhibited the sharpest divergence in affordance realisation, reflecting how deficiencies from earlier stages accumulated. Two affordance potentials were consistently observed: Local Agentic Acceleration (AP1) and Cognitive Integration & Coordination (AP3). AP1 in this stage depended heavily on orchestration and interoperability (AC-T2), literacy norms (AC-O1), and domain permissibility (AC-D1). In industry- and digital-native contexts, fragmented toolchains, manual handovers, and inconsistent documentation standards constrained agents to isolated operations such as generating draft code or updating single artefacts. AI-native firms, by contrast, possessed integrated repositories and orchestration frameworks that allowed agents to automate more coherent iteration routines - including documentation updates, test execution, and requirements reconstruction. AP3 became increasingly important during multi-team build/test cycles, but its realisation required literacy norms (AC-O1), semantic integration (AC-T1), orchestration (AC-T2), and domain conditions (AC-D1). Most organizations lacked this alignment, resulting in agents losing context across handovers or surfacing inconsistencies too late in the workflow. Develop thus demonstrates that advanced affordances materialise only when both upstream alignment and local integration foundations are in place.

Across all cases, the Deliver stage yielded no distinct stage-specific affordance potentials. Deliver work remains less digitised, more operationally embedded, and governed by strict process ownership, limiting both semantic integration (AC-T1) and orchestration (AC-T2). Agents contributed only marginally through isolated instances of Local Agentic Acceleration (e.g., summarising deployment documentation in AI-native settings), but these activities were residual outputs of earlier stages rather than true Deliver affordances. Instead, the only observed form of agentic support reflected Cross-Stage Continuity (AP5), where agents preserved context from upstream work. This absence of Deliver-specific affordances suggests that current agentic

capabilities align more naturally with exploratory and analytical knowledge work than with downstream implementation and scaling activities.

Cross-Stage Continuity (AP5) emerged as a widely needed but rarely realised affordance. All organizations experienced significant context loss between stages, which repeatedly led to rework, inconsistent assumptions, and cycles of misalignment. Partial realisation of AP5 appeared only within AI-native environments, where semantic integration (AC-T1), orchestration (AC-T2), literacy norms (AC-O1), evaluation capacity (AC-O2), governance clarity (AC-O3), and favourable domain conditions (AC-D1) were jointly present. In other settings, lacking one or more of these conditions, agents could not reliably maintain rationale, trace assumptions, or preserve historical decisions across stage boundaries. As the most demanding affordance - requiring the full set of technical, organizational, and domain conditions - AP5 remained largely unrealised across the sample. Its absence helps explain why innovation processes repeatedly "reset the logic" between stages, undermining continuity and diminishing the potential of earlier agentic contributions.

6.3 The Agentic Value Ladder

While the preceding sections analyse affordance potentials and their organizational boundaries within each stage of the innovation funnel, the cross-case pattern analysis revealed a set of dynamics that operated at a level above individual stages. These dynamics consistently shaped how value accumulated, stalled, or migrated across the funnel and could not be fully explained through stage-specific analysis alone. In line with abductive logic, this motivated a further step of theoretical integration, leading to the formulation of the Agentic Value Ladder. The model synthesizes how stage-level affordances interact systemically across the funnel and offers a complementary lens for understanding the broader value trajectory of agentic AI within organizational innovation processes.

Across the stage-level analysis in Section 6.2, a clear pattern emerges: organizations differ not only in which affordance potentials appear at each innovation stage, but more importantly in the configuration of actualization conditions that enable or constrain their realization. These patterns do not map neatly onto the funnel alone; instead, they reveal a deeper structural logic. Lower-order affordances tend to emerge where foundational technical conditions are in place, whereas higher-order affordances appear only in environments with more advanced organizational routines and governance. The cross-case comparisons also show that organizations often stabilize at different "levels" of realized value, reflecting distinct clusters of

AC sufficiency and insufficiency. To capture these cumulative and conditional relationships, the analysis synthesizes them into an agentic value ladder.

Figure 1 - The Agentic Value Ladder

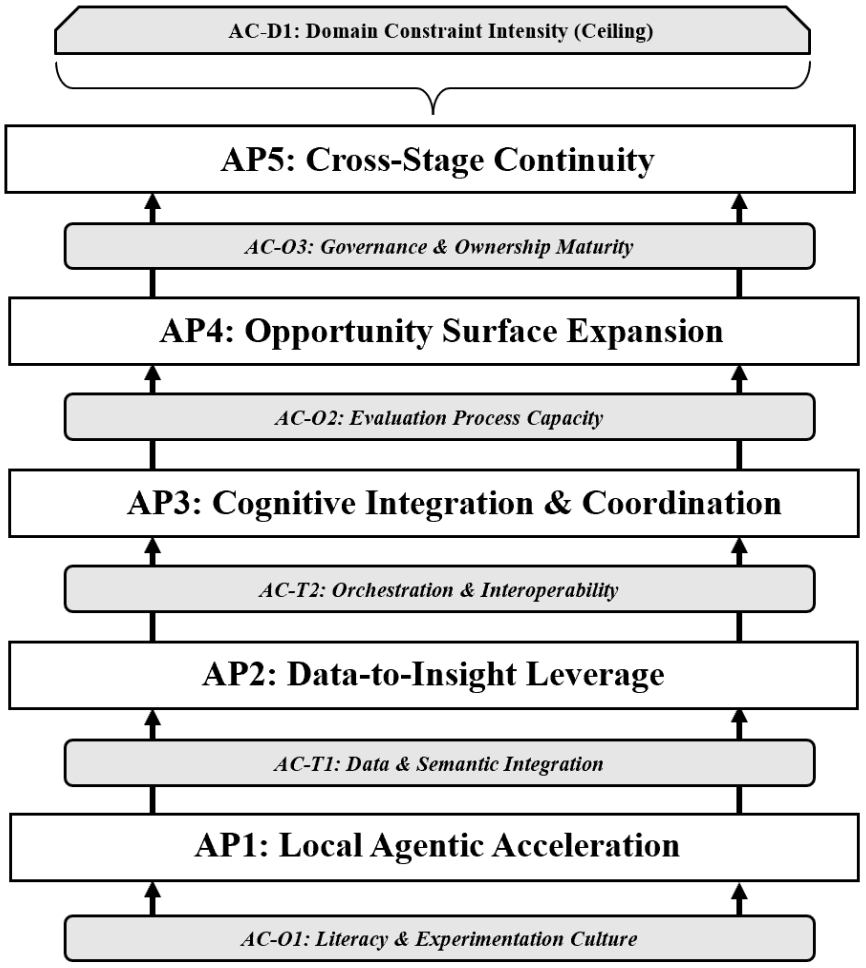


Figure 1 illustrates how affordance potentials (AP1-AP5) become accessible when specific actualization conditions (AC-T, AC-O, AC-D) are sufficiently developed. The ladder is not a linear maturity sequence but an ordered set of action possibilities. The empirics showed that higher-order affordances require more demanding combinations of technical and organizational foundations and are realized only when lower-order conditions are already in place.

At the base of the ladder, AP1 (Local Agentic Acceleration) is enabled when literacy and experimentation culture supports individual actors to test deployment of agents in their teams. AP2 (Data-to-Insight Leverage) requires semantic and data integration (AC-T1), which enables agents to synthesize distributed information. AP3 (Cognitive Integration and Coordination) depends on system orchestration and interoperability (AC-T2) so that agents can maintain shared

context across actors, tools, and workflows. AP4 (Opportunity Surface Expansion) relied on a stronger evaluation capacity (AC-O2), which allowed agents to broaden discovery beyond reactive sensing. At the top of the ladder, AP5 (Cross-Stage Continuity) depends on governance and ownership clarity (AC-O3). Beyond this all affordance potentials are bound by the organization's domain constraints, which both delimit value potential and create tensions accelerating AP realization, like the case of the eCommerce Digital-native's compliance pressure leading to the conceptualization of a compliance agent being planned for development.

The ladder is cumulative in that higher-order affordances rely on the stable realization of lower-order ones and cannot be sustained when foundational conditions remain weak. The interviews showed that organizations may demonstrate isolated instances of advanced affordances, but these do not scale without broader alignment in technical and organizational routines. Domain constraints (AC-D1) can also cap the upper bounds of actualization even when internal readiness is high.

Viewed through this structure, the ladder also clarifies why agentic maturity varies across organizations. Progression toward higher-order affordances requires not only technical readiness but also the orchestration of stable routines across teams, functions, and innovation stages. Organizations with fragmented workflows, inconsistent ownership, or weak cross-stage coordination are unable to sustain the conditions required for AP3-AP5, regardless of localized successes. Conversely, organizations with stronger orchestration capabilities are better positioned to achieve cumulative agentic value, even if their starting point is modest. In this sense, the ladder reflects the interplay between agentic capability and innovation-process maturity.

6.4 Mechanisms of the Agentic Value Ladder

Cross-case comparisons revealed three underpinning mechanisms that illuminate why organizations advance along the ladder in some contexts, stall in others, and at times regress. These mechanisms deepen the explanatory logic of the ladder and clarify the dynamic relationships between affordances and actualization conditions.

A first mechanism is bottleneck migration. As organizations begin to realize lower-order affordances such as Local Agentic Acceleration (AP1), new constraints emerge upstream or downstream in the innovation process. The empirical material showed that once acceleration became partially actualized, e.g. in local feasibility assessments or documentation tasks, attention shifted to previously latent bottlenecks in Data and Semantic Integration (AC-T1) or Orchestration and Interoperability (AC-T2). As these technical issues were addressed, organizational actualization conditions (AC-O2, AC-O3) surfaced, limiting progression toward

higher affordance potentials. Bottleneck migration explains why improvements at one rung of the ladder do not automatically enable ascent to the next; instead, new constraints continually arise as the locus of limitation shifts.

A second mechanism driving uneven progression on the agentic value ladder is the local-external actualization asymmetry. Several organizations demonstrated that even when internal conditions support localized agentic performance, progression stalls when downstream or external environments lack equivalent actualization conditions. This asymmetry was most visible in AI-native firms, where internal maturity frequently exceeded that of customer systems, data architectures, or governance routines, producing a ceiling on realized affordances regardless of internal capability. Industry-native and digital-native organizations displayed parallel patterns at smaller scales: pockets of internally feasible agentic augmentation could not translate into broader organizational value because adjacent teams, functions, or systems did not possess the necessary technical or organizational boundary conditions. This mechanism underscores that affordances are relational rather than intrinsic: an agent capable of high-order integration in one environment may only support low-order acceleration in another. Accordingly, AI-native firms progressed further up the ladder internally while remaining constrained in external deployments when surrounding environments failed to meet the minimum actualization conditions required for upward movement.

A third mechanism is the operational firefighting trap, which was particularly visible in the digital-native and, to some extent, the industry-native cases. In these environments, operational pressures and shifting priorities frequently displaced structured innovation work. As teams responded to immediate operational demands, experimentation norms deteriorated (AC-O1), evaluation capacities weakened (AC-O2), and continuity across stages broke down. This cycle prevented organizations from leveraging agents for higher-order affordance potential because the foundational organizational routines required to engage with agentic outputs were repeatedly interrupted. The firefighting trap thus represents a self-reinforcing mechanism that keeps organizations anchored on the lower rungs of the ladder.

Together, these mechanisms provide the dynamic explanatory logic of the agentic value ladder. They clarify how affordances accumulate, stall, or dissipate across stages; why progress is uneven or punctuated; and why organizations diverge despite common interest in agentic technologies.

7. Discussion

This chapter synthesizes the study's findings by linking the empirical results to the research questions and articulating their theoretical and practical implications. It is structured in four parts. First, we present the answers to the two research questions (7.1). Second, we outline the theoretical contributions of the study (7.2). Third, we consider the managerial implications (7.3). Finally, we discuss the study's limitations and identify opportunities for future research (7.4-7.5).

7.1 Conclusion

This thesis set out to explore how emerging agentic AI systems create value within innovation processes and how organizational settings shape their realisation. Existing research provides fragmented, task-level insights into how AI supports discrete innovation activities, but offers limited understanding of how AI reshapes interdependencies and coordination across the innovation funnel. At the same time, affordance–actualization research has focused mainly on micro-level, during-use interactions and has under-theorised the organizational and technical conditions that determine whether technological potentials can be actualized. Motivated by these gaps, we asked:

RQ1: *How can AI agents create value across the innovation funnel?*
RQ2*: How do organizational conditions enable or constrain the value potential of AI agents across the innovation funnel?*

In relation to RQ1, the findings show that agentic value manifests through five affordance potentials that appear at different levels of abstraction and that vary systematically by stage. Across Discover, Define, and Develop, these five affordance potentials manifest in stage-specific ways. The most immediately recognizable is *Local Agentic Acceleration* (AP1), which automates labour-intensive micro-tasks and delivers the efficiency gains managers typically associate with AI agents. However, the remaining affordances reflect capabilities that extend far beyond automation. *Data-to-Insight Leverage* (AP2) enables agents to synthesize fragmented information into coherent evaluative inputs; *Cognitive Integration and Coordination* (AP3) supports alignment of assumptions and interpretations across functions; and *Opportunity Surface Expansion* (AP4) broadens exploratory search to capture more opportunities and strengthen idea development and evaluation. A final cross-stage affordance, *Cross-Stage Continuity* (AP5), entails maintaining rationale, assumptions, and historical context as innovation work progresses from

Discover to Define and onwards to Develop. Notably, no stage-specific affordances materialised in Deliver, indicating that current agentic capabilities align more with upstream exploratory, digital and analytical work than with downstream operational implementation.

Turning to RQ2, this paper shows that affordances do not activate uniformly: each AP is bounded by a cumulative configuration of actualization conditions. Lower-order affordances such as AP1 require only organizational literacy & experimentation (AC-O1) and permissive domain constraints (AC-D1). Mid-level affordances (AP2 and AP3) depend critically on semantic integration (AC-T1), orchestration capabilities (AC-T2), and stable routines for interpretation and evaluation capacity (AC-O2). Higher-order affordances, especially AP5, demand the full alignment of conditions across all three domains: coordinated governance (AC-O3), semantic and orchestration maturity (AC-T1/T2), literacy and experimentation norms (AC-O1), evaluation culture (AC-O2), and permissive domain contexts (AC-D1). Across all cases, insufficient integration, weak cross-functional routines, and domain restrictions severely limited actualization of higher level affordance potentials beyond local acceleration, explaining the strong divergence between technical promise and organizational reality.

Integrating the answers to both research questions, the study proposes the Agentic Value Ladder as a mid-range theoretical contribution explaining how organizations progress from low-order to high-order agentic value. The ladder positions the five affordance potentials in an ordered sequence that reflects the increasing complexity of their enabling conditions. Lower levels require minimal literacy and opportunities for experimentation, while higher levels increase actualization conditions demands cumulatively. The mechanisms identified in the analysis, including bottleneck migration, local-external asymmetry, and the operational firefighting trap, account for why organizations often progress unevenly or stall despite strong interest in agentic technologies. Together, these contributions extend innovation-management research by clarifying how agentic AI can restructure the flow of innovation work and advance affordance–actualization theory toward autonomous, tool-integrating systems operating across multi-stage organizational processes.

7.2 Theoretical Contributions

This thesis offers three main contributions to affordance–actualization theory, innovation-process scholarship, and the emerging literature on agentic AI. The first contribution extends affordance theory by shifting analytical attention from micro-level, during-use enactment toward the organizational and technical preconditions that shape whether affordances can emerge at all. Prior research has predominantly focused on user-technology interactions and

reactive adjustments (Leonardi, 2011; Zammuto et al., 2007;), while repeatedly calling for organizational-level analyses of affordances and their structural determinants (Volkoff and Strong, 2017; Zammuto et al., 2007;). The AP–AC framework developed in this thesis answers this call by distinguishing affordance potentials as process-level opportunity structures and actualization conditions as the technical, organizational, and domain factors that govern their realizability. This moves affordance theory beyond its interactional origins and provides a proactive, multi-level account suited to technologies whose performance depends on autonomy, tool integration, and cross-stage reasoning rather than isolated features. In doing so, the model specifies not only what action possibilities agentic systems create, but what organizational foundations must be in place for those possibilities to materialise.

The second contribution advances innovation-process theory by showing how affordance dynamics vary systematically across the innovation funnel. Existing research describes the funnel as a sequence of activities involving exploration, evaluation, development, and implementation (Crossan and Apaydin, 2010; Hobday, 2005; Tidd and Bessant, 2018), but has not accounted for how emerging technologies introduce stage-contingent patterns of value creation. By mapping the interaction of affordance potentials and actualization conditions across Discover, Define, Develop, and Deliver, this thesis demonstrates that agentic value is shaped by the structural demands of each stage. Opportunity expansion is most prominent in discovery, cognitive integration becomes critical during definition and development, and acceleration spans all stages except Deliver. These stage-level dynamics are reinforced by three mechanisms identified in the analysis: bottleneck migration, local–external asymmetry, and the operational firefighting trap. Together, these mechanisms show how early-stage acceleration exposes downstream constraints, reframing the innovation funnel as an interdependent system in which upstream and downstream conditions co-evolve. This extends process theory by demonstrating that agentic AI reshapes not only individual tasks but the temporal pacing, sequencing logic, and cross-stage dependencies of innovation work.

The third contribution is to the emerging literature on agentic AI, where empirical theory-building has lagged behind technical development. While existing studies highlight the promise of autonomous tools, multi-step orchestration, and reasoning-driven workflows, most analyses remain anchored in model capabilities or architectural descriptions and offer limited insight into how organizations convert these capabilities into value. The Agentic Value Ladder developed in this thesis provides a middle-range theoretical model that explains how organizations progress from local acceleration toward systemic cognitive augmentation and cross-stage continuity. By linking each affordance potential to the corresponding configuration

of actualization conditions, the ladder clarifies why technological maturity alone does not predict agentic value and why organizations often stall at lower levels despite strong interest. This positions agentic AI not simply as a new class of digital artefact, but as a socio-technical system whose organizational value depends on cumulative alignment across technical infrastructure, organizational routines, governance, and domain constraints.

7.3 Managerial Implications

The managerial implications of this study should be interpreted as context-dependent and limited to the organizational environments represented in the cases. The findings do not claim universal generalizability; rather, they highlight patterns that may inform organizations operating under similar structural, technical, and domain conditions.

First, the AP–AC framework can support organizations in diagnosing where agentic AI is most likely to create value and what conditions currently constrain its realization. By clarifying which affordance potentials correspond to which enabling conditions, the framework helps identify where targeted investments in data integration, governance routines, or workflow design will have disproportionate impact.

Second, agentic value interacts with the innovation funnel in stage-contingent ways. Discovery benefits most from expanded search and early pattern synthesis. Define and Develop rely heavily on alignment, shared context, and structured evaluation routines. Deliver showed almost no stage-specific affordances in the studied contexts, and several cases illustrated that accelerating upstream work without evolving downstream processes generates bottlenecks. Within the limitations of our sample, organizations may use this pattern to sequence experimentation and anticipate where agentic value is likely to materialise first.

Third, the findings suggest that agentic AI creates system-level effects rather than isolated local gains. Improvements in feasibility checks, documentation, or analysis often exposed weaknesses elsewhere, particularly in data coherence, cross-team governance, and orchestration. Organizations therefore benefit from treating agent adoption as an organizational design question: strengthening coordination routines, shared data infrastructures, and cross-functional ownership structures is often a prerequisite for sustained value.

Fourth, adoption appears less as a discrete implementation and more as a progression. The Agentic Value Ladder provides a way for organizations to assess their current position, understand which actualization conditions are missing, and set realistic expectations for which forms of agentic value are attainable. The ladder emphasises that advancement depends on cumulative improvements rather than the deployment of new models or agents in isolation.

Lastly, the cross-case evidence indicates that organizational conditions are often more decisive than technical sophistication. Several technically advanced firms struggled to realise value due to fragmented workflows or unclear ownership, while others with modest tooling achieved gains through strong alignment, governance, and evaluation routines. Managerial attention may therefore be more productively directed at organizational readiness than at the technical frontier.

7.4 Limitations

The findings should be interpreted in light of several methodological limitations. Since the analysis operates at the organizational level while the data come from individual interview accounts, there is a risk that personal interpretations were extrapolated to organizational-level phenomena. This was mitigated through triangulation across senior leaders and functional experts, yet interpretive variation cannot be fully eliminated. The study further relies on self-reported descriptions of innovation processes and technical environments; confidentiality constraints limited access to internal systems or live agent deployments. As data were collected at a single point in time, the analysis reflects a cross-sectional snapshot rather than the longitudinal evolution of organizational practices and agent capabilities. Finally, although the cases span organizations with markedly different digital and AI maturities, all firms are Nordic. This geographic concentration may shape regulatory conditions, digital infrastructure, and innovation norms, and should therefore be considered when interpreting transferability.

Beyond methodological constraints, several theoretical and scope limitations also apply. The sample includes a small number of cases, chosen for theoretical rather than statistical representativeness; while suitable for theory-building, this restricts the breadth of empirical variation. The construction of affordance potentials, actualization conditions, and mechanisms is interpretive, meaning alternative theoretical lenses or coding choices may have yielded different structures. Moreover, the empirical focus was largely bounded to innovation processes; other organizational domains such as those with stricter regulatory oversight or operational constraints, may display different affordance patterns or boundary conditions. Taken together, the contributions of the study should be viewed as a plausible mid-range theoretical explanation of emerging agentic AI dynamics rather than a definitive or universally generalizable model.

7.5 Future Research

Future work could extend this study in several directions. Longitudinal research could examine how affordance potentials and enabling conditions evolve as organizations deepen adoption and as infrastructures mature. This would help determine whether the stage-dependent patterns observed here stabilize, shift, or diverge across industries. Further studies could broaden the empirical scope by exploring how agentic affordances manifest outside innovation processes, particularly in domains with high regulatory intensity or specialized workflows. Mixed-method or quantitative approaches could help test whether the mechanisms identified here generalize across larger samples and whether the Agentic Value Ladder can be operationalized into maturity assessments or comparative metrics. As agentic capabilities advance, research on governance, system interoperability, cross-team coordination, and domain constraints will be critical for understanding how organizations convert agentic potential into organizational value.

8. Bibliography

Aguirre, S. & Rodriguez, A., 2017. Automation of a business process using robotic process automation (RPA): A case study. *Applied Computer Sciences in Engineering. WEA 2017. Communications in Computer and Information Science*, 742. doi:10.1007/978-3-319-66963-2_7

Alavi, M. & Leidner, D.E., 2001. Knowledge management and knowledge management systems: Conceptual foundations and research issues. *MIS Quarterly*, 25(1), pp. 107–136. doi:10.2307/3250961

Anthropic, 2024. Model Context Protocol (MCP) technical specification. Available at: https://www.anthropic.com/research/model-context-protocol (Accessed: 7/12/2025).

Bailey, D.E. & Barley, S.R., 2020. Beyond design and use: How scholars should study intelligent technologies. *Information and Organization*, 30(2), 100286. doi:10.1016/j.infoandorg.2019.100286

Bao, Y., Gong, W. & Yang, K., 2023. A literature review of human–AI synergy in decision making: From the perspective of affordance actualization theory. *Systems*, 11(9), 442. doi:10.3390/systems11090442

Bell, E., Bryman, A. and Harley, B., 2022. *Business Research Methods*. Oxford University Press

Besson, P. & Rowe, F., 2012. Strategizing information systems-enabled organizational transformation: A transdisciplinary review and new directions. *Journal of Strategic Information Systems*, 21(2), pp. 103–124. doi:10.1016/j.jsis.2012.05.001

Bresnahan, T.F. & Trajtenberg, M., 1995. General purpose technologies: "Engines of growth"? *Journal of Econometrics*, 65(1), pp. 83–108. doi:10.1016/0304-4076(94)01598-T

Brynjolfsson, E., Li, D. & Raymond, L., 2025. Generative AI at work. *The Quarterly Journal of Economics*, 140(2), pp. 889–942. doi:10.1093/qje/qjae044

Brynjolfsson, E., & Mitchell, T., 2017. What can machine learning do? Workforce implications. *Science*, 358(6370), 1530-1534.

Brynjolfsson, E. & Hitt, L.M., 2000. Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4), pp. 23–48. doi:10.1257/jep.14.4.23

Chen, D., Doumeingts, G. & Vernadat, F., 2008. Architectures for enterprise integration and interoperability: Past, present and future. *Computers in Industry*, 59, pp. 647–659. doi:10.1016/j.compind.2007.12.016

Crossan, M.M. & Apaydin, M., 2010. A multi-dimensional framework of organizational innovation: A systematic review of the literature. *Journal of Management Studies*. doi:10.1111/j.1467-6486.2009.00880.x

Dealroom, 2025. Opening moves in global AI venture capital: $110 billion invested in AI-first companies in 2024. *Dealroom Global AI VC Report*. Available at: https://dealroom.co/uploaded/2025/02/AI-Summit-2025.pdf (Accessed: 7/12/2025).

Design Council, 2005. The double diamond. Available at: https://www.designcouncil.org.uk/our-resources/the-double-diamond/ (Accessed: 7/12/2025).

Du, W.Y., Pan, S.-L., Xie, K. & Xiao, J., 2020. Data analytics contributes to better decision-making beyond organizational boundaries. *MIS Quarterly Executive*, 19(2), Article 5. Available at: https://aisel.aisnet.org/misqe/vol19/iss2/5 (Accessed: 7/12/2025).

Dunphy, S. M., Herbig, P. R., & Howes, M. E. (1996). The innovation funnel. *Technological Forecasting and Social Change*, 53(3), 279–292. doi:10.1016/S0040-1625(96)00098-4

Enholm, I.M., Papagiannidis, E., Mikalef, P. & Krogstie, J., 2022. Artificial intelligence and business value: A literature review. *Information Systems Frontiers*, 24, pp. 1709–1734. doi:10.1007/s10796-021-10186-w

Eisenhardt, K. M. (2021). What is the Eisenhardt Method, really? *Strategic Organization*, 19(1), 147–160. doi: 10.1177/1476127020982866

Eloundou, T., Manning, S., Mishkin, P. & Rock, D., 2023. GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*. doi:10.48550/arXiv.2303.10130

European Commission, 2024. Artificial Intelligence in the European Commission (AI@EC). Brussels: European Commission. Available at: https://commission.europa.eu/system/files/2024-01/EN%20Artificial%20Intelligence%20in%20the%20European%20Commission.PDF (Accessed: 7/12/2025).

Fayard, A.L. & Weeks, J., 2014. Affordances for practice. *Information & Organization*, 24, pp. 236–249. doi:10.1016/j.infoandorg.2014.10.001

Financial Times, 2025. How much of the market is AI? *Financial Times*. Available at: https://www.ft.com/content/1edc4113-4282-474d-928b-4b84b7f2044e (Accessed: 7/12/2025).

Foong, S.B., Tan, K.S. and Lee, Y.K. (2013) 'An insight into an approach to IT architecture: A best practice example, and its social aspects, from a leading technology company', *Procedia Technology*, 11, pp. 441–449. doi: 10.1016/j.protcy.2013.12.323.

Füller, J., Tekic, Z. & Hutter, K., 2024. Rethinking innovation management - how AI is changing the way we innovate. *Journal of Business Research*, 60(4). doi:10.1177/00218863241287323

Gama, F. & Magistretti, S., 2023. Artificial intelligence in innovation management: A review of innovation capabilities and a taxonomy of AI applications. *Journal of Product Innovation Management*, pp. 1–36. doi:10.1111/jpim.12698

Geertz, C. (1973) *The Interpretation of Cultures*. New York: Basic Books.

Gibson, J.J., 1977. The theory of affordances. In: R. Shaw & J. Bransford, eds. *Perceiving, acting, and knowing*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 127–137.

Gioia, D.A., Corley, K.G. & Hamilton, A.L., 2013. Seeking qualitative rigor in inductive research. *Organizational Research Methods*, 16(1), pp. 15–31. doi:10.1177/1094428112452151

Glaser, B.G. & Strauss, A.L., 1967. *The discovery of grounded theory: Strategies for qualitative research*. Chicago: Aldine.

Godé, C. & Brion, S., 2024. The affordance-actualization process of predictive analytics: Towards a configurational framework of a predictive policing system. *Technological Forecasting and Social Change*, 204, 123452. doi:10.1016/j.techfore.2024.123452

Händler, T., 2023. A taxonomy for autonomous LLM-powered multi-agent architectures. In: *Proceedings of the 15th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2023)*. doi:10.5220/0012239100003598

Haefner, N., Wincent, J., Parida, V. & Gassmann, O., 2021. Artificial intelligence and innovation management: A review, framework, and research agenda. *Technological Forecasting and Social Change*, 162, 120392. doi:10.1016/j.techfore.2020.120392

Hoessler, S., 2024. Guiding incumbent companies in navigating digital transformations: A qualitative study on structural ambidexterity and strategic leadership. *Journal of Entrepreneurship, Management and Innovation*, 20(4), pp. 49–72. doi:10.7341/20242043

Humberd, B.K. & Latham, S.F., 2025. When AI becomes an agent of the firm: Examining the evolution of AI in organizations through an agency theory lens. *Journal of Management Studies*. doi:10.1111/joms.13274

Hussmo, D., Säfsten, K. & Wlazlak, P., 2022. Engineering tools as boundary objects between product development and production. In: *Transdisciplinarity and the future of engineering*. IOS Press, pp. 687–696. doi:10.3233/ATDE220702

Hutchby, I., 2001. Technologies, texts and affordances. *Sociology*, 35(2), pp. 441–456. doi:10.1177/S0038038501000219

IBM, 2025. Orchestrating agentic AI for intelligent business operations. *IBM Institute for Business Value*. Available at: https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/agentic-process-automation (Accessed: 7/12/2025)

International Organization for Standardization (ISO), 2022. *ISO/IEC 22989: Information technology – Artificial intelligence – Artificial intelligence concepts and terminology*. Geneva: ISO. Available at: https://www.iso.org/standard/74296.html (Accessed: 7/12/2025)

Jöhnk, J., Weißert, M. & Wyrtki, K., 2021. Ready or not, AI comes: An interview study of organizational AI readiness factors. *Business & Information Systems Engineering*, 63(1), pp. 5–20. doi:10.1007/s12599-020-00676-7

Jordan, M.I. & Mitchell, T.M., 2015. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), pp. 255–260. doi:10.1126/science.aaa8415

Keller, R., Stohr, A., Fridgen, G., Lockl, J. & Rieger, A., 2019. Affordance-Experimentation-Actualization theory in artificial intelligence research – a predictive maintenance story. In: *Proceedings of the 40th International Conference on Information Systems (ICIS 2019)*

Lepak, D.P., Smith, K.G. & Taylor, M.S., 2007. Value creation and value capture: A multilevel perspective. *Academy of Management Review*, 32(1), pp. 180–194. doi:10.5465/amr.2007.23464011

Leonardi, P.M., 2011. When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. *MIS Quarterly*, 35(1), pp. 147–167. doi:10.2307/23043493

Lincoln, Y.S. & Guba, E.G., 1985. *Naturalistic inquiry*. Newbury Park, CA: Sage Publications.

Mariani, M. & Dwivedi, Y.K., 2024. Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research*, 175, 114542. doi:10.1016/j.jbusres.2024.114542

Markus, M.L. & Silver, M.S., 2008. A foundation for the study of IT effects: A new look at DeSanctis and Poole's concepts of structural features and spirit. *Journal of the Association for Information Systems*, 9, doi:10.17705/1jais.00176.

Microsoft, 2024. Introduction to Copilot Agents. *Microsoft Learn*. Available at: https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-overview (Accessed: 7/12/2025).

Nisa, U. et al., 2025. agentic AI: The age of reasoning - A review. *Journal of Automation and Intelligence*. doi:10.1016/j.jai.2025.08.003

Oberländer, A., Roeglinger, M. & Rosemann, M., 2021. Digital opportunities for incumbents – A resource-centric perspective. *The Journal of Strategic Information Systems*, 30, 101670. doi:10.1016/j.jsis.2021.101670

Orlikowski, W.J. & Scott, S.V., 2008. Sociomateriality: Challenging the separation of technology, work and organization. *The Academy of Management Annals*, 2(1), pp. 433–474. doi:10.1080/19416520802211644

Park, J.S., O'Brien, J.C., Cai, C.J., Morris, M.R., Liang, P. & Bernstein, M.S., 2023. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*. doi:10.48550/arXiv.2304.03442

Patton, M.Q., 2015 *Qualitative Research & Evaluation Methods: Integrating Theory and Practice*. 4th edn. Thousand Oaks, CA: SAGE Publications.

Roberts, D.L. & Candi, M., 2024. Artificial intelligence and innovation management: Charting the evolving landscape. *Technovation*, 136, 103081. doi:10.1016/j.technovation.2024.103081

Russell, S. & Norvig, P., 2021. *Artificial intelligence: A modern approach*. 4th ed. Harlow: Pearson.

Schumpeter, J.A., 1934. *The theory of economic development*. Cambridge, MA: Harvard University Press.

Sequoia Capital, 2025. Our companies. Available at: https://sequoiacap.com/our-companies/#all-panel (Accessed: 7/12/2025).

Shrestha, Y.R., Ben-Menahem, S.M. & von Krogh, G., 2019. Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), pp. 66–83. doi:10.1177/0008125619862257

Soori, M., Arezoo, B. & Dastres, R., 2023. Artificial intelligence, machine learning and deep learning in advanced robotics: A review. *Cognitive Robotics*, 3, pp. 54–70. doi:10.1016/j.cogr.2023.04.001

Strauss, A. & Corbin, J., 1998. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. 2nd ed. Thousand Oaks, CA: Sage.

Strong, D.M. & Volkoff, O., 2009. Understanding organization–enterprise system fit: A path to theorizing the information technology artifact. *MIS Quarterly*, 34(4), pp. 731–756.

Tech Monitor, 2025. Global AI venture capital hits US$110bn in 2024 - driven by foundational models. *TechMonitor*. Available at: https://www.techmonitor.ai/digital-economy/ai-and-automation/global-ai-venture-capital-110bn-2024-driven-foundational-models (Accessed: 7/12/2025).

Teece, D.J., 2007. Explicating dynamic capabilities: The nature and microfoundations of sustainable enterprise performance. *Strategic Management Journal*, 28(13), pp. 1319–1350. doi:10.1002/smj.640

Tekic, Z. & Füller, J., 2023. Managing innovation in the era of AI. *Technology in Society*, 73, 102254. doi:10.1016/j.techsoc.2023.102254

The Economist, 2024. Just how rich are businesses getting in the AI gold rush. *The Economist*. Available at: https://www.economist.com/business/2024/03/17/just-how-rich-are-businesses-getting-in-the-ai-gold-rush (Accessed: 7/12/2025)

Tidd, J. & Bessant, J., 2018. *Managing innovation: Integrating technological, market and organizational change*. 6th ed. Chichester: Wiley.

Timmermans, S. & Tavory, I., 2012. Theory construction in qualitative research: From grounded theory to abductive analysis. *Sociological Theory*, 30(3), pp. 167–186. doi:10.1177/0735275112457914

Truong, Y. & Papagiannidis, S., 2022. Artificial intelligence as an enabler for innovation: A review and future research agenda. *Technological Forecasting and Social Change*, 183, 121852. doi:10.1016/j.techfore.2022.121852

Vernadat, F.B., 2007. Interoperable enterprise systems: Principles, concepts, and methods. *Annual Reviews in Control*, 31(1), pp. 137–145. doi:10.1016/j.arcontrol.2007.03.004

Volkoff, O. & Strong, D.M., 2017. Affordance theory and how to use it in IS research. In: R. Galliers & M.-K. Stein, eds. *The Routledge companion to management information systems*. London: Routledge, pp. 232–245.

von Krogh, G., 2018. Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Academy of Management Discoveries*, 4(4), pp. 404–409. doi:10.5465/amd.2018.0084

Weill, P. & Ross, J.W., 2004. *IT governance: How top performers manage IT decision rights for superior results*. Boston, MA: Harvard Business School Press.

White House OSTP, 2023. Blueprint for an AI Bill of Rights: Implementation Guidance. Washington, DC: OSTP. Available at: https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/ (Accessed: 7/12/2025).

Yin, R.K., 2003. *Case study research: Design and methods*. 3rd ed. Thousand Oaks, CA: Sage Publications.

Zammuto, R.F., Griffith, T.L., Majchrzak, A., Dougherty, D.J. & Faraj, S., 2007. Information technology and the fabric of organization. *Organization Science*, 18(5), pp. 749–762. doi:10.1287/orsc.1070.0307

Appendices

Appendix A - Semi-Structured Interview Protocol

Section 1. Project Background & Conceptual Understanding

Purpose	Example Prompts
Introduce the project and clarify objectives.	• Confirm that the respondent has read the preparatory document. • Briefly restate the purpose of the study: to explore the value potential of AI agents across the innovation funnel and understand enabling or constraining conditions for their implementation.

Section 2. Respondent Background

Purpose	Example Prompts
Understand the respondent's background, role, and responsibilities.	• What is your background and current role? • If you were to break down the main responsibilities of your department/role into different parts or functions, what would those be? Roughly what share of your time or focus goes to each (adding up to 100%)? • Does your work follow any recurring cycle or sequence - such as planning, execution, evaluation, and implementation? If so, what are those cycles?

Section 3. Core Questions on Innovation and AI

Theme	Example Prompts
Innovation Contribution	• In what ways does your work contribute to innovation within the organization? • How is innovation structured or supported within your department or function?
Stages of Involvement	• How do you identify or contribute to new opportunities? • Do you help define or refine potential solutions? • Do you test new approaches or tools before deciding to scale or implement them fully? • Have you been involved in implementing or scaling something new?
Organizational Enablers & Barriers	• Are there any factors in your organization that make it easier - or harder - for new ideas to emerge or be implemented? • Have you implemented AI before? What have been factors of success and failure?

AI agent Potential	• What opportunities for value creation do you see for AI agents in your innovation process? • What challenges or constraints do you foresee?

Section 4. Optional Technical Module (for technical respondents)

Used selectively when interviewing participants with architectural or technical expertise. In interviews with nontechnical respondents the interview proceeded to deep dive into the role and emerging themes of that specific respondent and interview.

Theme	Goal	Example Prompts
Enterprise Architecture Overview	Understand the firm's technical foundation.	• How would you describe your Enterprise Architecture across infrastructure, application, data, and integration layers? • What are the main technical challenges or enablers of AI and AI agents in your current setup?
Technology Infrastructure	Identify cloud posture and compute environment.	• Are your core workloads primarily on-prem, in the cloud, or hybrid? • Which cloud providers do you currently use?
Application Environment	Map where AI agents could operate.	• What are your main business systems (ERP, CRM, PLM, ticketing, etc.)? • Are these mostly legacy, SaaS, or internally built platforms? • Do these systems expose APIs or SDKs for integration? • Do employees already interact with AI or automation tools?
Information Architecture	Assess data accessibility and structure.	• Where is operational data stored (databases, data lake, warehouse)? • Are there shared or standardized data models? • How is data access governed (role-based, ad-hoc)?
Integration Layer	Understand system connectivity and data exchange.	• What tools or platforms do you use for system-to-system integration (API gateways, iPaaS, event buses)? • Do you support event-driven or real-time data exchange, or mostly batch processes?

Appendix B - Respondents Summary

Case	Department	Job Title	Focus
Industry-native #1	Product Development	Head of Product Development	Product Development

Industry-native #1	Business Development	Market Developer	Business Development & Systems
Industry-native #1	IT	IT System Lead	AI, IT Architecture & Technical Conditions
Industry-native #2	Management	CEO	Full Scope
Industry-native #2	Management	CTO	Product & Innovation
Digital-native #1	Strategy & Business Development	Strategy & Business Development Manager	Cross-functional Innovation & Internal Processes
Digital-native #1	Tech	Chief Technology Officer	AI, IT Architecture & Technical Conditions
Digital-native #1	Merchandising	Supply Intelligence & Automation Manager	Previous AI Implementations
Digital-native #1	Commercial	Chief Commercial Officer	AI & Pull Innovation
Digital-native #1	Merchandising	Supply Intelligence & Automation Manager	Merchandising AI Affordances
Digital-native #1	Product	Chief Product Officer	Product & Innovation
Digital-native #2	Board of Directors	Head of the Board	Full Scope
Digital-native #2	Management	Chief Technology Officer	Technical Scope
AI-native #1	Management	CEO	Full Scope
AI-native #1	Customer Success	Customer Success Lead	Customer's AI agent potential
AI-native #2	Management	Technical Co-founder	Full Scope
AI-native #2	Management	Co-founder	Full Scope
AI-native #3	Management	Co-founder & CEO	Full Scope
AI-native #3	Management	Technical Co-founder	Full Scope

Appendix C - Case Category Patterns and Quotes

Case Category	Pattern	Quote(s)
Industry-native	Fragmented data architecture and missing semantic integration	*"We have some problems with the registers not always being harmonized… the same attribute can be called 1072 in one system and something entirely different in another. It becomes difficult to match without building a translation register [...] we often end up piecing things together from old files, emails, and scattered systems because nothing is linked all the way through."* *- Market Developer, Industry-native #1*
	Resource-intensive innovation processes with overloaded evaluation capacity	*"When we work with these really large [pre-studies] you have to calculate everything manually, request supplier quotes, summarize everything, put it all in one big Excel file and calculate the final cost [...] even the smallest feasibility check means going system to system, checking specs, calling suppliers, verifying documents… it always takes longer than you think."* *- Head of Product Development, Industry-native #1* *"There's never enough time for proper evaluations. The urgent things always win, so pre-studies pile up."* *- IT System Lead, Industry-native #1*
	Scattered nature of innovation discovery	*"New business opportunities can be driven by regulatory requirements, shortages of raw materials, or that we want to start producing a product ourselves. We get ideas that are customer-driven or from ourselves […] we look at import statistics, who produces the product today, which customers buy it, the volumes, and whether it fits our terminal strategy."* *- Market Developer, Industry-native #1*
	High perceived potential for agentic AI across the innovation processes	*"[AI agent systems] is really interesting… One agent monitors customer and technology trends, another generates new ideas, a third analyzes business value and proposes pilots, and a fourth documents and spreads learning. It becomes like a chain."* *- Head of Product Development, Industry-native #1* *"If we could just harmonize the data, a lot of the scanning and initial analysis could be automated. That's where agents could save a massive amount of time."* *- Market Developer, Industry-native #1* *"At our inception, we thought the capabilities of our [hardware] was the bottleneck to our business. After launching our first commercial product, we discovered that the massive amount of data that was spit out of our [product] was the real bottleneck. Making sense of all that data required loads of manhours but revealed insights that proved instrumental for our product development - we could see what issues we needed to focus on but also where we created the most value to our customers. Now we could deploy [AI solution provider] to do the same thing in seconds."* *- CEO, Industry-native #2*
	Presence of a supportive organizational culture but insufficient structural readiness for agent adoption	*"We built [the company] from nothing with limited resources … our priority was to build a disruptive technology and lead a fundamental shift in the sensor technology industry, not creating the best business operating system. [Therefore], our organization was designed to drive innovation within our hardware and software-on-chip solutions - governance and digitalization was secondary. Now, as a mature publicly listed company, we have shifted more attention to address this digitalization debt and see a lot of potential in AI across our departments."* *- CEO, Industry-native #2* *"The culture is positive, but IT resources are limited. We can't maintain too many new solutions without clearer governance."* *- IT System Lead, Industry-native #1* *"People are curious and want to try new things, but the structure around digital projects isn't strong enough. It becomes unclear who should own what."* *- Market Developer, Industry-native #1*

	Safety, compliance, and domain-specific constraints requiring human judgment	*"We operate in both life science and automotive, where non-compliance can kill the business, so we've always been extremely careful with new software and have likely overinvested in manual quality and safety checks. At the same time, our core mission is to improve safety and accessibility, and many of our biggest early breakthroughs came from adopting new digital solutions."* *- CTO, Industry-native #2* *"You can't rely on automation when the consequences of an error are so high. Someone with experience always has to make the final call."* *- Head of Product Development, Industry-native #1*
Digital-native	Data and insight infrastructure constraints	*"We have so much data that it's really hard to dig into it and understand it. Besides, no one really knows how to read the data properly… so a lot of insight work becomes manual guesswork and right now we're not utilizing [the data] enough."* *- Chief Product Officer, Digital-native #1* *"The data is there, but there's no process in place to actually make sense of it or use it systematically."* *- Strategy & Business Development Manager, Digital-native #1* *"Frontend [E-commerce] operations rely on scattered app integrations and data sources to build [state-of-the-art] features that optimize [sales] conversion […] we lack the ability to handle the complexity of our own [technical] setup or make use of our own analytics data [...] so we rely heavily on external agencies to improve the customer experience."* *- Chief Technology Officer, Digital-native #2* *People think AI means 'ask ChatGPT to write some copy'. They have no idea how much context and structure is required to automate real storefront operations. We don't have that structure today."* *- Chief Technology Officer, Digital-native #2*
	Innovation process immaturity and structural weaknesses	*"Prioritization is a problem… it's very top-down and the process isn't really structured…Sometimes a leader brings up an idea and then we evaluate it, but from the bottom-up side it's very ad-hoc."* *- Strategy & Business Development Manager, Digital-native #1* *"Most of what's considered important comes from leadership… we don't really have a clear structure for how ideas should move forward."* *- Brand & Loyalty Manager, Digital-native #1* *"If we want to refresh our user experience on the storefront, we usually end up overpaying consultants or see underwhelming results from an increasingly overloaded tech team. Yet, making the wrong UX changes carry massive revenue impact so we know we are missing out if we don't find a way to iterate faster "* *- Head of the Board, Digital-native #2*
	Operational firefighting systematically limiting organizational capacity for innovation efforts	*"We have more ideas than tech capacity […] so it's always hard prioritization. We've had a peak of fires because of recent platform migration - everything else has had to wait until after Black Week."* *- Chief Technology Officer, Digital-native #1* *"It's difficult to get cross-functional time… innovation always loses when operations are under pressure."* *- Brand & Loyalty Manager, Digital-native #1* *"Every time the platform shakes, all innovation stops. We have no buffer. Teams are so stretched that we survive quarter to quarter. You can forget systematic experimentation or AI adoption under those conditions."* *- Head of the Board, Digital-native #2*
	Fragmented technological environments and missing orchestration	*"It's extremely scattered with very different knowledge levels… and we haven't started with any orchestration yet."* *- Supply Intelligence & Automation Manager, Digital-native #1* *"We don't want to get locked in, so we have many different tools and no platform that ties it all together."* *- Supply Intelligence & Automation Manager, Digital-native #1* *"Technical experimentation is difficult because the systems don't talk to each other."* *- Chief Product Officer, Digital-native #1*
	Emerging AI adoption and capability building	*"We had an AI week [and] many people went from thinking AI was just prompting to understanding that there are agents, workflows, vibe-coding and much more. Many reported that [LLM application] was very good for bouncing ideas while others used it for repetitive tasks like documentation and long e-mails."* *- Brand & Loyalty Manager, Digital-native #1*

		"There's a lot of interest in agents for daily routines, but adoption is very uneven depending on the team." *- Supply Intelligence & Automation Manager, Digital-native #1*
	Leadership signals and cultural dynamics	*"The CEO is head-over-heels when it comes to AI… but the level of [technical] understanding across the organization is very uneven."* *- Strategy & Business Development Manager, Digital-native #1* *"Most of what's considered important comes from leadership."* *- Brand & Loyalty Manager, Digital-native #1* *"Learning doesn't spread equally… some teams experiment a lot while others barely get started."* *- Supply Intelligence & Automation Manager, Digital-native #1*
AI-native	Agentic automation restructuring internal product development velocity	*"It took two days to fix a UX in the front-end, now it takes five minutes."* *- CEO, AI-native #1* *"[Our AI agent operating system] is dynamic, multimodal, and responsive in real-world environments. Guests see instant results on-screen maps, menus, bookings, or stay details, while the agent can also call a restaurant, notify a shuttle driver, or trigger a maintenance task."* - Co-founder, *AI-native #2*
	Realization of agentic innovation depends on guided onboarding, education, and cultural readiness	*"In the beginning people thought it was scary… but now it's completely natural. It will be the same with AI agents. People need handholding at first, because hiring an AI agent is still new."* *- Customer Success Lead, AI-native #1* *"We're not asking hotels to change behavior, we're automating existing workflows with better technology. [Our solution] isn't a rip-and-replace. It's middleware. That means low IT lift, fast deployment, and less risk."* *- Technical co-founder, AI-native #2*
	Technical integration foundations determine the ceiling of agentic affordances	*"When we onboard a customer, we go into Slack or Teams and then connect GitHub and their ticket-management system… and then they're up and running."* *- Customer Success Lead, AI-native #1* *"Across [agentic AI] use cases, it usually comes down to context - finding the right context and enabling in-context reasoning. This means connecting to [the customer's] preferred apps and systems and extracting the relevant information. The hardest part for the agent isn't executing the task, it's understanding why it's being done and how it fits into everything else."* *- Technical co-founder, AI-native #3* *"Our MCP architecture allows real-time data sync and action execution that these platforms can't match without significant re-architecture. They would need to rebuild from scratch to offer our level of integration."* *- Co-founder, AI-native #2*
	Staged maturity path of agentic AI realization in innovation processes	*"Implementation usually happens in three phases… first AI search and extracting value from existing data, then handling data silos, and only later the proper workflows constructed across tools."* *- CEO, AI-native #1*
	Agentic systems surface hidden innovation opportunities by unifying fragmented internal and external data	*- CEO, AI-native #1* *"The core value of our product comes from helping people manage complexity, especially when digital information, communication, and decisions are spread across many channels and the information load makes it hard to navigate between what's important and what's simply noice"* *- Co-founder & CEO, AI-native #3*
	Market readiness and organizational maturity limits AI agent adoption for customers	*"Human-in-the-loop will be crucial… People don't want a black box making decisions; they want a partner."* *- CEO & co-founder, AI-native #3* *"There's huge potential when the only thing missing is that people don't understand yet… adoption depends on the ecosystem they already have in place."* *- CEO, AI-native #1*
	Agents act as cross-team cognitive integrators by aligning divergent mental models	*"People are often aligned on what they think is going well, but not aligned on what they think is going badly… AI agents can map everyone's mental models and pull out the real core of the problem so the team can solve it together."* *- Customer Success Lead, AI-native #1*

"Across problems, what matters is understanding why a task is being done and how it fits into everything else… that's where current tools fall short and where an AI agent can truly make a difference."
- Technical co-founder, AI-native #3
"An agent that remembers in a structured way but can still improvise and reason beyond predefined paths helps people coordinate around the actual context rather than assumptions."
- Co-founder & CEO, AI-native #3

Appendix D - AI Disclosure

AI tools employed	The thesis made use of ChatGPT 5.1 and Gemini 3.
Contribution to thesis quality	The tools improved the speed and efficiency at which we could ensure linguistic and structural quality in the thesis by supporting rephrasing, smoothing transitions, and ensuring consistent academic tone across sections. In addition, they could serve as a "third party" for critique or review of ideas or material to ensure triangulation at a frequency not possible just through formal supervision or peer-review. Most importantly, the tools enabled a more coherent narrative and higher-quality analytical synthesis by shifting researcher time from mechanical tasks to conceptual reasoning.
Risks encountered and mitigation measures	Key risks included the possibility of stylistic over-reliance that could reduce authorial integrity. To mitigate this risk, all text refined by AI underwent manual review to ensure theoretical accuracy and consistency with our analytical intent. No unverified AI-generated material was included.

Insights gained from the process	The use of ChatGPT 5.1 and Gemini 3 demonstrated that AI tools are particularly effective in accelerating high-volume mechanical tasks such as re-writing undesirable sentences or identifying non-academic sources. Their value was highest when used iteratively as partners rather than generative substitutes. The process highlighted the necessity of rigorous human verification to maintain methodological reliability while showing that AI can meaningfully enhance depth, consistency, and throughput in academic work when integrated thoughtfully into the research workflow.