

Enterprise Cognitive Product Strategy Engineering for AI-Driven Retail Supply Chains

Ashok Bandi

Independent Researcher, USA. ORCID: 0009-0007-5442-1424

Abstract: Retail supply chains now operate cognitive systems capable of processing millions of demand signals per hour, yet the product strategy frameworks governing those systems were designed for human-paced decision cycles. This misalignment is not merely a speed problem; it is an architectural one. When AI systems outrun the governance models meant to oversee them, enterprises are forced into a binary that undermines the value of their investments: constrain cognitive agents to tasks narrow enough for human oversight cycles, or accept autonomous operation without adequate governance. This paper proposes the Cognitive Product Strategy Engineering (CPSE) Stack, a six-layer operational architecture that resolves this tension by embedding AI capability directly into the strategic product decision-making process. The six layers - Strategic Intent Formation, Enterprise Signal Capture, Cognitive Reasoning Engine, Product Orchestration, Adaptive Governance, and Outcome Learning - form an integrated architecture grounded in design science research methodology. The framework is evaluated through three scenario-based validation cases representing AI maturity progression in retail supply chain contexts. A five-level, five-dimension AI maturity assessment instrument accompanies the framework to enable practitioners to self-locate within the CPSE development pathway. Results confirm that the CPSE Stack provides a structured, governance-ready architecture for advancing enterprise AI product strategy beyond experimentation toward institutionalized intelligence. The framework is intended for researchers studying cognitive product systems, practitioners scaling AI in retail supply chains, and organizations seeking alignment between AI capability and strategic accountability.

Keywords: Enterprise cognitive product strategy engineering; AI-driven supply chain; cognitive product strategy; retail supply chain intelligence; product orchestration; decision intelligence; AI governance; maturity assessment; design science research

1. INTRODUCTION

Retail supply chains now operate cognitive systems capable of processing millions of demand signals per hour, yet the product strategy frameworks governing those systems were designed for human-paced decision cycles. This misalignment is not merely a speed problem; it is an architectural one.

The specific structural failure is this: existing product strategy frameworks treat AI as a capability to be incorporated into human decision workflows rather than as a system that must be governed through engineered architecture. When AI inference speed, data volume, and interconnection density exceed the cognitive bandwidth available to product and supply chain teams, the result is not AI-augmented human decisions but AI decisions that outrun governance entirely.

This paper positions enterprise cognitive product strategy engineering as the discipline required to close that gap. The central question it addresses is not whether AI should govern supply chain product decisions - current deployment trajectories make that a near-certainty for large-scale retail - but how product strategy frameworks should be architected to govern AI-driven supply chains without sacrificing either cognitive speed or accountability.

Three gaps in the existing literature motivate this contribution. First, the scholarly literature on AI in retail supply chains [2][3] has matured to cover recommendation systems, demand forecasting, and autonomous fulfillment, but it does not address the product strategy governance layer that sits above those operational systems. Second, product strategy frameworks in management information systems have been updated to accommodate AI co-pilots and



decision-support tools [4][6], but these frameworks still assume human decision-makers at the center of strategic choices; they do not address the governance architecture required when AI systems execute strategic intent autonomously. Third, responsible AI frameworks such as the NIST AI Risk Management Framework [1] provide governance principles but do not prescribe product strategy architecture for supply chain AI specifically.

This paper addresses all three gaps through a single contribution: the Cognitive Product Strategy Engineering (CPSE) Stack, a six-layer operational architecture that integrates AI capability, enterprise signal processing, cognitive reasoning, product orchestration, adaptive governance, and outcome learning into a coherent engineering framework for retail supply chain product strategy.

The paper is organized as follows. Section 2 reviews the literature across five relevant streams. Section 3 presents the research methodology. Section 4 describes the CPSE Stack architecture. Section 5 presents scenario-based validation across three archetype cases. Section 6 introduces the AI maturity assessment instrument. Section 7 discusses theoretical and practical implications. Section 8 concludes.

2. LITERATURE REVIEW

2.1 Agentic AI in Retail Supply Chain Management

The deployment of autonomous AI agents in supply chain contexts has accelerated significantly over the past three years. Syed et al. [2] demonstrate that multi-agent architectures applied to supply chain optimization yield approximately 30% stockout reduction and 15% holding cost improvement relative to heuristic baselines - confirmed metrics that establish the operational case for cognitive supply chain systems. Zheng et al. [3] extend this foundation, showing that cognitive signal processing at enterprise scale enables demand pattern recognition across product portfolios that exceeds human analytical capacity by orders of magnitude. The consistent finding across this literature is that agentic AI in retail supply chains is not a future capability but a present deployment reality, and the governance question is therefore urgent.

Microsoft's industry analysis [15] confirms the practical adoption of agentic supply chain systems, documenting retail deployments in which AI agents coordinate across inventory positioning, demand sensing, and supplier communication functions. The documented implementations share a common architectural characteristic: the AI systems operate within defined parameters but execute autonomously within those parameters, without human oversight at individual decision points. This operational reality is what makes product strategy governance architecture a research priority rather than a forward-looking speculation.

2.2 Product Strategy Frameworks and AI Integration

Trumpler et al. [4] examine how enterprises are adapting product strategy processes to incorporate AI capabilities, identifying a consistent pattern: AI tools are added to existing planning architectures designed for quarterly review cycles - frameworks built around the cognitive bandwidth, decision authority, and temporal rhythm of human planning teams. Their finding that this integration model produces suboptimal outcomes is directly relevant to the present paper's motivation. Guo and Mantravadi [6] extend this work by proposing integration patterns for generative AI within supply chain decision frameworks, demonstrating that AI adoption in supply chain management has advanced significantly, with enterprises moving from exploratory tools to operational deployment at scale. The adoption predictor literature within this stream [17] identifies technology fit - whether system capability actually maps onto the decision types practitioners must navigate - as the primary predictor of realized AI value. The CPSE Stack is designed with this finding as a central engineering constraint.

2.3 Agentic Systems and Multi-Agent Coordination

Jannelli et al. [5] make a strong case for the feasibility of multi-agent consensus-seeking to disrupt the entire flow of supply chain, showing that specialists functioning in parallel, collaborating with demand response, inventory positioning and supplier orchestration, achieve consensus and cooperate towards common goals. Their findings - about

30% fewer stockouts and about 15% fewer holding costs over stockout baselines - demonstrate the compounding gains from multi-agent coordination. D. Ivanov [7] advance this line through agentic digital twins: architectures that bridge model-based simulation with AI-driven autonomous decision-making, representing the integration of continuous reasoning with high-fidelity environment modeling. These two works - together with Mohammed et al. [8], who document AI-augmented product orchestration in retail environments - establish the technical foundation on which the CPSE Stack's Product Orchestration layer is built.

2.4.2.4 Cognitive Intelligence in Product Design - Historical and Contemporary Perspectives

Grigorev et al. [18] document the evolution of cognitive intelligence in product design from early computational approaches through contemporary AI integration, establishing the intellectual lineage within which the CPSE Stack operates. Researchers extend this to manufacturing systems, showing how cognitive architectures can be embedded in product design processes at scale [14]. Taken together, these works position the CPSE Stack not as a novel departure but as the extension of an established cognitive design tradition to the specific domain of product strategy governance - a dimension that existing maturity models do not assess.

2.5 Responsible AI Governance Frameworks

The NIST AI Risk Management Framework [1] provides the most operationally actionable structure for enterprise AI governance, organizing governance responsibilities under four functions: Govern, Map, Measure, and Manage. Its emphasis on continuous risk assessment positions governance as an ongoing operational activity rather than a one-time policy exercise - an orientation that directly informs the CPSE Stack's Adaptive Governance layer. Gothi [10] addresses autonomy-accountability boundaries in AI-assisted product decisions, establishing that AI governance must be treated as an architectural design problem - determining where to draw the line between machine-executable action and human-declared intent - rather than as a policy imposed on top of a running system. Gothi's framework requires explicit autonomy level definitions at each operational layer, a principle the CPSE Stack operationalizes throughout its architecture. To counter the hallucination problem in a cognitive reasoning system, Sreeram [11] shows that retrieval-augmented generation (RAG) architectures offer groundedness mechanisms for enterprise product knowledge.

2.6 Generative AI and Large Language Model Applications in Supply Chain

An example of using structured Knowledge Reps to enable reliable reasoning in centralized Enterprise AI deployments for SC decision support is demonstrated by Sun, Wan et al. [16]. Aslam, Saleem, and Lai [17] present a systematic review of the current applications of generative AI and ChatGPT for supply chain management, concluding that there is ample proof of efficiencies gained with AI and ChatGPT, but technology fit is emerging as a reason for not implementing the technology. Bandara et al. [9] document a production-grade agentic AI framework for retail supply chain operations, showing that human-in-the-loop orchestration preserves accountability while automating decision-intensive workflows at scale. Becklines and El-Gayar [12] establish precedent for applying design science research to AI system architectures for operational decisions, providing the methodological warrant for the present paper's framework development contribution.

2.7 Gap Synthesis

Across these five streams, a consistent gap emerges: the literature addresses AI systems (agentic supply chain tools), AI governance principles (NIST AI RMF, AAB model), and AI adoption patterns (technology fit, generative AI applications), but does not address product strategy as an engineering discipline that integrates these three domains into a single operational architecture. Cognitive product design research [14][18] addresses processes and outcomes; but the layer connecting them - product strategy as an engineering discipline - is not addressed in either body of work. The CPSE Stack fills this gap.

3. RESEARCH METHODOLOGY

This paper adopts a design science research (DSR) methodology [12], following the established tradition of producing artifacts - frameworks, models, instruments - that address identified organizational problems. DSR is appropriate here for three reasons: the problem is an organizational and architectural one (how to govern AI-driven product strategy in retail supply chains) rather than an empirical one; the contribution is a framework rather than a statistical finding; and scenario-based validation is the appropriate evaluation approach for frameworks that operate at the architectural level.

The research design follows the DSR cycle in five phases: problem identification (structural misalignment between AI capability and product strategy frameworks); objective definition (governance-ready architecture for AI-driven product strategy); artifact design (the CPSE Stack and maturity instrument); artifact evaluation (scenario-based validation across three archetypes); and communication (this paper). Becklines and El-Gayar [12] have established precedent as a legitimate contribution form - specifically in framework papers addressing governance and decision architecture questions in AI-enabled organizations.

The article is intentionally written in generalized industry language and avoids identifying any employer-specific project, confidential implementation, internal system, or proprietary workflow. The framework is designed to be applicable across large-scale retail supply chain environments and is evaluated against generalized scenario archetypes rather than against specific organizational data. Vallé et al. [13] confirm that cognitive architecture frameworks for industrial AI strategy developed through DSR methodology produce contributions that are both theoretically grounded and practically deployable, validating this methodological choice for the present work.

4. THE COGNITIVE PRODUCT STRATEGY ENGINEERING (CPSE) STACK


4.1 Architectural Overview

The CPSE Stack organizes enterprise AI product strategy into six sequential but tightly coupled layers. The central architectural principle is that governance belongs inside the stack, not on top of it - a design response to Gothi's [10] finding that governance imposed as a policy overlay on an autonomous system cannot adapt at the speed the system operates. Each layer has defined inputs, outputs, governance triggers, and failure modes. Cross-layer feedback loops operate continuously, enabling the system to recalibrate in response to both market signals and governance constraints.

Table 1: CPSE Stack - Layer Summary

Layer	Name	Primary Function	Key Components	Governance Trigger
L1	Strategic Intent Formation	Translate enterprise objectives into machine-executable intent specifications	OKR decomposition engine; AI mandate definition; boundary conditions	Annual strategic review; material business event (market entry, regulatory change)
L2	Enterprise Signal Capture	Aggregate and synthesize real-time market, inventory, and demand signals at scale	IoT integration; POS data streams; supplier feeds; external market data	Signal anomaly threshold breach; data quality degradation alert
L3	Cognitive Reasoning Engine	Convert synthesized signals into product decisions through AI reasoning and multi-agent coordination	RAG-grounded LLM reasoning; multi-agent consensus protocols;	Confidence score below threshold; novel decision type; regulatory sensitivity flag

			threshold-governed escalation	
L4	Product Orchestration	Execute AI-generated product decisions across functional domains	Cross-functional workflow automation; supplier communication agents; inventory repositioning	Transaction value threshold; cross-regional scope; contractual commitment level
L5	Adaptive Governance	Ensure AI actions remain within strategic intent, ethical boundaries, and regulatory requirements	NIST AI RMF functions (Govern, Map, Measure, Manage); AAB model; audit trail	Autonomy boundary approach; audit event; regulatory change notification
L6	Outcome Learning	Close the strategic loop by feeding decision outcomes back to L1 for intent recalibration	Performance attribution; decision-outcome linkage; strategy recalibration triggers	Performance variance exceeding tolerance; strategic drift detection



[FIGURE 1: CPSE Stack Architecture Diagram]

4.2 Layer 1 - Strategic Intent Formation

Strategic intent formation is the top layer of the CPSE Stack and the one that preserves human strategic accountability in an architecture that otherwise delegates extensively to cognitive systems. Its function is to translate enterprise objectives - revenue targets, service level commitments, sustainability mandates, competitive positioning goals - into machine-executable intent specifications: precise, bounded, and assessable statements of what the AI systems below are authorized to optimize for and within what constraints.

The key design principle of L1 is specificity over generality. Strategic intent expressed as 'improve supply chain performance' cannot be reliably operationalized by AI systems, which require unambiguous optimization targets and boundary conditions. Intent expressed as a hierarchy of weighted objectives with explicit constraint sets is machine-executable and auditable. Strategic intent is reviewed annually and whenever a material business event requires strategic reconfiguration.

4.3 Layer 2 - Enterprise Signal Capture

Enterprise signal capture addresses one of the most consequential limitations of conventional product strategy processes: their dependence on aggregated, periodically refreshed data rather than real-time signal streams. In retail supply chains, the relevant signals - point-of-sale velocity, inventory positions, supplier capacity, weather events, competitive pricing movements, logistics disruptions - arrive continuously and interact in ways that static data snapshots cannot represent. Microsoft's analysis of agentic retail deployments [15] confirms that real-time signal synthesis is achievable at enterprise scale and is foundational to the performance gains documented in agentic supply chain systems.

L2 functions as the sensory layer of the CPSE Stack. Its outputs are not raw data streams but synthesized cognitive state representations: compressed, semantically coherent summaries of enterprise signal state that L3 can reason over without requiring access to underlying data infrastructure. The distinction matters for governance: by establishing L2 as a defined architectural layer with its own governance triggers, the CPSE Stack creates an auditable record of what the AI reasoned from, not merely what it decided.

4.4 Layer 3 - Cognitive Reasoning Engine

Layer 3 receives the enterprise cognitive state from L2 and applies AI reasoning to generate product decision recommendations and, within defined autonomy thresholds, product decisions. A multi-agent architecture - with specialist agents operating in parallel across demand response, inventory positioning, and supplier orchestration - is the recommended implementation pattern, drawing on the multi-agent consensus-seeking work of Jannelli et al. [5] and the knowledge graph-enhanced reasoning demonstrated by Sun, Wan et al. [16].

Critically, L3 incorporates retrieval-augmented generation grounding [11] as an architectural requirement, not an optional enhancement. The distinction between a memorized pattern and a grounded fact is not academic - it is the difference between a recommendation that reflects current inventory positions and a recommendation that reflects what current inventory positions typically look like. Sreeram's [11] RAG grounding mechanisms directly address this hallucination risk in enterprise product knowledge contexts. If L3 needs to escalate a decision to human reviewers because the threshold for decisions that can be autonomously made is exceeded, L5's governance protocols are used, enforcing the Autonomy-Accountability Boundary model [10] at the reasoning layer in addition to after the fact.

4.5 Layer 4 - Product Orchestration

The product orchestration process takes the AI-driven product decisions from L3 and turns them into action across functional areas: procurement, inventory management, logistics, supplier management, and retail execution. Mohammed et al. [8] document the operational patterns of AI-augmented product orchestration in retail environments, establishing that cross-functional automation is the primary mechanism through which agentic supply chain systems deliver their documented efficiency gains.

The governance design of L4 reflects the insight that autonomy thresholds must be operationally specific. A transaction value threshold that triggers human approval is more precise and more auditable than a general directive to 'escalate high-stakes decisions.' L4 maintains a complete action log that feeds L6's outcome attribution function, enabling the system to learn which orchestration patterns produce which outcomes across different market conditions. The knowledge graph architecture described by Sun, Wan et al. [16] provides the data structure that makes this attribution tractable at enterprise scale.

4.6 Layer 5 - Adaptive Governance

Adaptive governance is the most distinctive layer of the CPSE Stack and the one that most directly addresses the structural gap identified in the literature review. Existing AI governance frameworks - including the NIST AI RMF [1] - provide comprehensive principles but do not address the architectural question of where governance logic should reside within an operating AI system. The CPSE Stack's answer is that governance belongs inside the stack, not outside it.

L5 implements the four core functions of the NIST AI RMF - Govern, Map, Measure, and Manage - as continuous operational processes rather than periodic assessment activities. Its connection to the Autonomy-Accountability Boundary (AAB) model [10] is explicit: L5's function is to maintain the AAB as a live operational boundary rather than as a static policy document. Governance that cannot adapt to the system it governs is not governance - it is a compliance checkbox with a latency problem. L5's adaptive mechanisms address this by monitoring autonomy boundary proximity and triggering governance review before boundaries are breached rather than after.

4.7 Layer 6 - Outcome Learning

Outcome learning closes the strategic loop by feeding decision outcomes back to L1 for intent recalibration. If not, cognitive product strategy systems could be good on deployment day, but never get better - they could actually do worse as market conditions, strategic priorities and supply chain environments change over time or as a result of the decisions the systems make on a daily basis about the market.

In cognitive retail platforms, Bandara et al. [9] report on a deployed multi-agent retail framework, which illustrates that supply-chain-scale automation with preserved human oversight is operationally feasible: The feedback loop of translated decisions and strategic retuning allows for measurable continuous improvement. The agentic digital twin architecture of D. Ivanov [7] provides the simulation infrastructure that makes outcome learning robust: by modeling the counterfactual outcomes of decisions not taken alongside the actual outcomes of decisions executed, L6 can distinguish between effective strategies and strategies that appeared effective due to favorable market conditions.

5. SCENARIO-BASED VALIDATION

Three scenario archetypes were constructed to evaluate the CPSE Stack across a representative range of retail supply chain product strategy challenges. Each archetype represents a distinct strategic decision type - time-sensitive demand response, network optimization, and supply resilience - and each places distinct demands on different layers of the stack, enabling a cross-sectional evaluation of architectural completeness.

5.1 Scenario Archetype 1 - Demand Volatility Response

When there is a sudden surge of 40pc demand in certain product stocks in hundreds of stores across the region, the inventory positions are already low due to previous decisions regarding supply chain planning stocked by the retailer, and at the same time, demand is being further intensified by competitor stock-outs. Without the CPSE Stack, the event escalates through a human decision chain - category manager, regional director, supply chain operations - typically producing a generic response four to eight hours after the triggering event.

With the CPSE Stack active, L2 detects the demand signal pattern and synthesizes a cognitive state update within minutes of the event onset. L3's demand response agent evaluates available inventory positions, supplier

capacity, and logistics feasibility, generating a ranked set of response actions with projected stockout probability and supply chain feasibility - a level of specificity that a human decision chain cannot produce at the same speed. L4 executes authorized responses autonomously, escalating only the subset of decisions - large cross-regional inventory transfers - that exceed the authorized threshold and require human approval through L5's governance protocol. L6 records outcome data from the event, enabling the system to improve its response pattern for future occurrences [2].

5.2 Scenario Archetype 2 - Inventory Positioning at Scale

Rebalancing excess inventory across a national store network - thousands of nodes, hundreds of thousands of active SKUs, competing objectives - is the defining optimization challenge of large-scale retail supply chain management. The CPSE Stack's multi-agent architecture, drawing on the coordination mechanisms of Jannelli et al. [5], enables this problem to be addressed as a continuous optimization exercise rather than a periodic planning exercise.

L3's multi-agent coordination framework evaluates proposed transfers for cascade effects - detecting whether resolving one excess position would create a shortage elsewhere - using the agentic digital twin infrastructure described by D. Ivanov [7]. The result is that inventory positioning decisions that previously required multi-day planning cycles can be generated, evaluated, and executed within hours, transforming a periodic planning exercise into a continuous product management capability - compressing the gap between optimal strategy and executed action from days to hours.

5.3 Scenario Archetype 3 - Supplier Orchestration Under Disruption

A tier-1 supplier declares a production constraint affecting 15% of active product lines, requiring rapid qualification of alternative suppliers, renegotiation of existing agreements, and update of supply chain systems with the resulting arrangements - a process measured in days in conventional supply chain operations and constrained by relationship complexity and information asymmetry.

L3's supplier orchestration agent - applying the multi-agent consensus-seeking mechanisms of Jannelli et al. [5] - evaluates alternative suppliers against a weighted scoring framework derived from the enterprise's strategic intent specification in L1. L5's governance protocol ensures that final contract terms are escalated to human approval, maintaining the Autonomy-Accountability Boundary at the commitment point. L6 records supplier performance data from the disruption event, feeding the qualification model used by L3 in future events.

Table 2: Scenario Validation Matrix

Dimension	Scenario A Demand Volatility	Scenario B Inventory Positioning	Scenario C Supplier Orchestration
AI Maturity Level	Level 2-3 (Defined to Integrated)	Level 3-4 (Integrated to Managed)	Level 3-4 (Integrated to Managed)
Stack Layers Active	L1, L2, L3 (demand agent), L4, L5	L1, L2, L3 (inventory agents), L4, L5, L6	L1, L2, L3 (supplier agent), L4, L5, L6
Key Decision Type	Demand response prioritization and inventory release	Multi-node inventory optimization	Supplier qualification and contract coordination
Governance Trigger	Cross-regional inventory transfer threshold	Large single-site transfer; cascade-risk flag	Contract commitment; tier-1 supplier change
Human Override Point	Cross-regional transfer approval	Transfers above defined value threshold	Final contract term approval
Key References	[2], [8], [10]	[5], [6], [7]	[5], [10], [11]

6. AI MATURITY ASSESSMENT INSTRUMENT

6.1 Instrument Design and Five Dimensions

The AI maturity assessment instrument introduced in this section provides practitioners with a structured mechanism for self-locating within the CPSE development pathway and identifying priority investment areas. Five dimensions correspond to a cluster of CPSE Stack capabilities.

Dimension 1 - Strategic Alignment: The degree to which enterprise AI investments are governed by explicit, machine-executable strategic intent specifications rather than ad hoc business case justification.

Dimension 2 - Signal Integration: The completeness and real-time fidelity of enterprise signal capture: the extent to which the organization has unified siloed data sources into a continuous cognitive state representation.

Dimension 3 - Cognitive Autonomy Governance: The sophistication of the autonomy-accountability boundary framework: whether the organization has explicit, enforced autonomy thresholds at each layer, or operates with general governance statements and threshold-governed autonomous action [10].

Dimension 4 - Orchestration Maturity: The scope and speed of AI-driven cross-functional action: whether AI decisions are executed in near real-time across functional domains or require multi-step human routing before execution becomes autonomous within defined boundaries.

Dimension 5 - Outcome Learning Velocity: The speed and granularity of outcome-to-strategy feedback: the time elapsed between a decision's execution and its performance data being incorporated into strategic intent recalibration (L6 and L1).

6.2 Five-Level Maturity Rubric

Level 1 - Ad Hoc: No systematic process. AI tools are used inconsistently. Strategic intent is informal. Signal data is fragmented, and outcome attribution is anecdotal.

Level 2 - Defined: Processes are documented but unevenly followed. Strategic intent exists in planning documents but is not machine-executable. AI governance is asserted through policy, not architecture.

Level 3 - Integrated: All CPSE Stack layers are deployed and integrated. Signal capture is real-time. L3 operates with defined autonomy thresholds. Governance is architectural. Outcome learning operates on defined cycles.

Level 4 - Managed: Maturity is quantitatively tracked. Governance boundaries are dynamically adjusted based on system performance. Outcome learning velocities strategy recalibration on accelerated cycles.

Level 5 - Autonomous and Self-Optimizing: The CPSE Stack operates as a self-governing system within enterprise-declared strategic boundaries. Strategy recalibration is continuous. Human oversight is reserved for boundary redefinition.

Figure 2: CPSE Maturity Radar




FIGURE 2: CPSE Maturity Radar - Visualization

The three validation scenarios from Section 5 map onto the maturity radar to illustrate the instrument's practical application. Scenario A (demand volatility response) requires high maturity on Dimensions 1, 3, and 4; it can tolerate lower maturity on Dimension 5 because the learning cycle is secondary to the immediate response capability. Scenario B (inventory positioning) requires the highest maturity on Dimensions 2 and 4, as the optimization quality is directly dependent on signal fidelity and orchestration reach. Scenario C (supplier orchestration) requires the highest maturity on Dimension 3, as the governance sophistication required for supplier commitment decisions is the primary capability constraint.

A consistent finding across all three scenarios is that enterprises investing in reasoning capability (L3 - models and multi-agent architectures) without commensurate investment in governance architecture (L5 - AAB implementation, NIST AI RMF functions) expose themselves to correlated failure modes - locally optimized decisions that create systemic fragility [1][10]. The maturity instrument is designed to surface this imbalance explicitly, flagging Dimension 3 gaps when Dimension 4 maturity is high.

7. DISCUSSION

7.1 Theoretical Contributions

The CPSE Stack contributes to three bodies of scholarly literature. To the cognitive product strategy literature [13][14], it contributes an engineering architecture that treats product strategy as a designed system rather than a planning process - extending the cognitive computing tradition to the governance dimension [18]. To the supply chain AI literature [2][5][7], it contributes a governance-native framework that specifies where autonomy thresholds belong in a layered architecture and how outcome learning connects to strategic intent recalibration. To the responsible AI literature [1][10], it provides a concrete illustration of how NIST AI RMF functions can be operationalized as architectural layers rather than governance overlays.

The paper extends the DSR methodology precedent established by Becklines and El-Gayar [12] to the specific domain of product strategy governance - a dimension that existing maturity models do not assess. The five-dimension maturity instrument contributes a self-assessment mechanism that bridges the gap between framework adoption and measurable maturity progression.

7.2 Practical Implications

Product managers operating in retail supply chain contexts face an increasingly urgent architectural decision: whether to govern AI systems through policy overlays on existing product strategy processes or to architect AI

capability directly into the product strategy system. The CPSE Stack argues, both theoretically and through scenario validation, that policy overlay is architecturally insufficient for systems operating at cognitive speeds.

The practical implication of the maturity instrument is diagnostic before it is prescriptive. Technology fit - confirmed as the primary predictor of adoption success [17] - is achieved by design in the CPSE Stack through layer-level capability matching. A practitioner who identifies low maturity on Dimension 3 (Cognitive Autonomy Governance) while operating high Dimension 4 (Orchestration Maturity) has a specific, actionable risk profile rather than a generic AI capability gap. Microsoft's [15] documentation of enterprise AI deployment patterns in retail confirms that this diagnosis - strong orchestration without adequate governance architecture - is the most common failure mode in large-scale retail AI deployments.

7.3 Governance Implications

The alignment between the CPSE Stack and the NIST AI RMF [1] is intentional and operationally specific. L5's implementation of the Govern, Map, Measure, and Manage functions as continuous architectural processes rather than periodic reviews directly addresses the AI RMF's emphasis on dynamic risk management. The AAB model [10] provides the boundary-definition framework that L5 requires but does not itself specify: where autonomy ends and human accountability begins at each layer and for each decision type.

7.4 Limitations and Future Research

The CPSE Stack is evaluated through scenario archetypes rather than empirical deployment data, which limits the strength of causal claims about performance outcomes. The framework's applicability to retail supply chain contexts with lower data infrastructure maturity has not been addressed. The interaction between the CPSE Stack's Layer 2 signal capture requirements and the privacy-preservation constraints of emerging data governance regulations requires dedicated investigation.

Future research directions include: empirical validation of the maturity instrument across diverse retail supply chain contexts; integration of the CPSE Stack with supplier-side cognitive systems for end-to-end supply chain intelligence; and exploration of the framework's applicability beyond retail - to healthcare logistics, manufacturing, and financial services - to test its generalizability across supply chain-intensive industries [6][13].

8. CONCLUSION

The structural misalignment between AI capability and product strategy governance in retail supply chains is not a temporary condition that will resolve as enterprises gain AI experience. It is an architectural problem that requires an architectural solution: a framework that embeds governance into the product strategy system rather than imposing it on top.

The Cognitive Product Strategy Engineering (CPSE) Stack proposed in this paper provides that architecture: six layers - Strategic Intent Formation, Enterprise Signal Capture, Cognitive Reasoning Engine, Product Orchestration, Adaptive Governance, and Outcome Learning - integrated into a coherent engineering framework that treats product strategy as a designed system. The scenario validation demonstrates that the stack is architecturally complete across the three core retail supply chain product strategy challenges: demand volatility response, inventory positioning at scale, and supplier orchestration under disruption.

The AI maturity assessment instrument introduced alongside the stack provides practitioners with a diagnostic mechanism calibrated to the specific governance and capability requirements that large-scale retail AI deployment demands. Its core argument - that cognitive product strategy is an engineering problem, not a governance overlay - has direct implications for how enterprises approach AI investment, organizational design, and strategic accountability in retail supply chain contexts.

The CPSE Stack's most significant theoretical contribution is the operationalization of governance as an architectural design problem rather than a policy question. Future work should validate this argument empirically and extend the framework to supply chain contexts beyond retail.

REFERENCES

1. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, Gaithersburg, MD, Jan. 2023. [Online]. Available: <https://doi.org/10.6028/NIST.AI.100-1>
2. T. A. Syed, S. Jan, G. Ali, A. Akarma, A. Ali, and Q.-A. Mastoi, "Agentic AI Framework for Smart Inventory Replenishment," arXiv preprint arXiv:2511.23366, Nov. 2025. [Online]. Available: <https://arxiv.org/abs/2511.23366>
3. L. Zheng, L. Yang, Z. Li, G. Lyu, C. M. Koh, and C.-P. Teo, "Adapting to Evolving Requirements: Agentic AI for Retail Supply Chain Operations," arXiv preprint arXiv:2609.03860, Sep. 2026. [Online]. Available: <https://arxiv.org/abs/2609.03860>
4. L. Trumpler, R. F. de Assis, E. R. da Silva, L. A. de Santa-Eulalia, and C. Hendriksen, "Agentic AI Autonomy Assessment: A Decision-Support Framework Towards Governed Supply Chain Systems," arXiv preprint arXiv:2607.25405, Jul. 2026. [Online]. Available: <https://arxiv.org/abs/2607.25405>
5. V. Jannelli, S. Schoepf, M. Bickel, T. Netland, and A. Brintrup, "Agentic LLMs in the Supply Chain: Towards Autonomous Multi-Agent Consensus-Seeking," 2024, DOI: 10.1080/00207543.2025.2604311. [Online]. Available: <https://arxiv.org/abs/2411.10184>
6. D. Guo and S. Mantravadi, "The role of digital twins in lean supply chain management: review and research directions," *International Journal of Production Research*, vol. 63, no. 5, 2025, DOI: 10.1080/00207543.2024.2372655. [Online]. Available: <https://doi.org/10.1080/00207543.2024.2372655>
7. D. Ivanov, "Agentic digital twins: bridging model-based and AI-driven decision-making support for a new era of supply chain and operations management," *International Journal of Production Research*, vol. 64, no. 15, pp. 6574-6590, 2026, DOI: 10.1080/00207543.2026.2630277. [Online]. Available: <https://doi.org/10.1080/00207543.2026.2630277>
8. I. A. Mohammed, S. Verma, J. Srinivasan, H. P. Josyula, S. R. K. Sudhamathi, and A. D. Todmal, "AI-Driven Supply Chains: Enhancing Efficiency and Sustainability in Modern Commerce," *International Journal of Computer Information Systems and Industrial Management Applications*, vol. 18, no. 19s, pp. 819-833, 2026. [Online]. Available: <https://csupub-ijcिसim.org/index.php/ijcिसim/article/view/5085>
9. E. Bandara et al., "Flowr — Scaling Up Retail Supply Chain Operations Through Agentic AI in Large Scale Supermarket Chains," arXiv preprint arXiv:2604.05987, Apr. 2026. [Online]. Available: <https://arxiv.org/abs/2604.05987>
10. E. Bandara et al., "Think Before You Act: A Neurocognitive Governance Model for Autonomous AI Agents," arXiv preprint arXiv:2604.25684, Apr. 2026. [Online]. Available: <https://arxiv.org/abs/2604.25684>
11. B. Sreeram, "Retrieval-Augmented Generation: Architecting Trustworthy Knowledge Systems on Top of Large Language Models," *International Journal of Computer Information Systems and Industrial Management Applications*, vol. 18, no. 20s, pp. 710-720, 2026. [Online]. Available: <https://csupub-ijcिसim.org/index.php/ijcिसim/article/view/5213>
12. L. Becklines and O. El-Gayar, "Dimensions of Artificial Intelligence Maturity Models in Supply Chain Management: A Systematic Literature Review," in *Proc. Americas Conference on Information Systems (AMCIS) 2025*, AIS Electronic Library. [Online]. Available: https://aisel.aisnet.org/amcis2025/sig_osra/sig_osra/14/
13. R. Valle, J. Seitz, D. K. Kanbach, and M. C. Schuhmacher, "The use of artificial intelligence in new product development: A systematic literature review, conceptual framework, and future research agenda," *Technovation*, vol. 153, art. no. 103541, 2026, DOI: 10.1016/j.technovation.2026.103541. [Online]. Available: <https://doi.org/10.1016/j.technovation.2026.103541>
14. "Towards cognitive intelligence-enabled product design: The evolution, state-of-the-art, and future of AI-enabled product design," *Journal of Manufacturing Systems*, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S2452414X24002024>
15. F. Miller, O. Guy, and P. Manikas, "Enhancing supply chain efficiency in the retail and consumer goods industry with agentic systems," *Microsoft Cloud Blog*, Feb. 13, 2025. [Online]. Available: <https://www.microsoft.com/en-us/microsoft-cloud/blog/retail-and-consumer-goods/2025/02/13/enhancing-supply-chain-efficiency-in-the-retail-and-consumer-goods-industry-with-agentic-systems/>
16. Y. Sun, L. Wan et al., "Uncovering the Interactions Between Enterprise AI Transformation, Supply Chain Concentration, and Corporate Risk-Taking Capacity," *IEEE Transactions on Engineering Management*, vol. 71, 2024, DOI: 10.1109/TEM.2024.3411631. [Online]. Available: <https://ieeexplore.ieee.org/document/10555306/>
17. J. Aslam, A. Saleem, and K.-H. Lai, "Supply Chain Management in the Era of Generative AI (ChatGPT): Technology Fit and Psychological Drivers of Adoption," *IEEE Transactions on Engineering Management*, 3817- 3831, vol. 72, 2025, DOI: 10.1109/TEM.2025.3605823. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/11151815/>
18. "The Evolution, Framework, and Future of Cognitive Intelligence-enabled Product Design," *Procedia CIRP*, vol. 109, 2022, DOI: 10.1016/j.procir.2022.05.238. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2212827122007387>