Augmentation or overwhelm? GenAI and the recoding of supply chain planning's DNA

Citation for the original published paper (version of record):
Shurrab, H., Jonsson, P. (2026). Augmentation or overwhelm? GenAI and the recoding of supply chain planning's DNA. International Journal of Physical Distribution and Logistics Management, 56(11): 241-264. http://dx.doi.org/10.1108/IJPDLM-08-2025-0444

Augmentation or overwhelm? GenAI and the recoding of supply chain planning's DNA

International Journal of Physical Distribution & Logistics Management

Hafez Shurrab
Department of Management, College of Business Administration, Ajman University, Ajman, United Arab Emirates, and
Patrik Jonsson
Department of Technology Management and Economics, Chalmers University of Technology, Gothenburg, Sweden

Received 24 August 2025
Revised 5 January 2026
24 March 2026
Accepted 1 April 2026

Abstract

Purpose – Supply chain management literature describes generative AI (GenAI) as transformative for operations, but its socio-technical consequences for the professional workforce remain underexplored. This study investigates how GenAI adoption reshapes core supply chain planning (SCP) roles.

Design/methodology/approach – Employing an exploratory multi-case study design, the study compares job specifications from GenAI adopter firms (Amazon, Tesla, Colgate-Palmolive and The Warehouse Group) with those of matched non-adopter firms across three deployment architectures. A strict separation between classification data (strategic documents, executive statements and technical publications) and analysis data (job specifications) prevents circular reasoning. Semi-structured interviews with senior SCP leaders were triangulated with the textual analysis to reveal day-to-day practices that formal documentation does not capture.

Findings – Two different archetypes emerge: the process guardian, who executes procedures within transaction-focused systems and the supply chain architect, who orchestrates adaptive planning across AI-enabled platforms. GenAI adoption produces an autonomy–ambiguity paradox, whereby planner authority expands while the decision space becomes harder to define. Formal hiring documentation lags behind operational deployment across firms. Four transition-specific paradoxes characterize the progression from early to advanced GenAI maturity in SCP roles.

Originality/value – A transformation framework models pathways from GenAI deployment to augmentation or overwhelm. A three-category typology of deployment maturity (GenAI-native, GenAI-augmented and build-phase) captures variations that binary adopter/non-adopter classification would collapse. A maturity model operationalizes this framework through diagnostic stages that comprise transition paradoxes and resolution requirements. Nine propositions structure future research on human–AI collaboration in SCP.

Keywords Generative artificial intelligence, Supply chain planning, Human-AI collaboration, Workforce transformation, Maturity model, Case study

Paper type Research article

1. Introduction

Technologies that automate existing tasks offer efficiency but often create new cognitive demands for human operators (Raisch and Krakowski, 2021). The nature of those demands varies by the mode of AI. Generative AI (GenAI) [1] shifts human labor from content production to content curation and verification (Chalmers *et al.*, 2026). GenAI deployments in supply chain planning (SCP), such as Amazon's transformer-based forecasting models (Ansari *et al.*, 2024) and Colgate-Palmolive's implementation of large language models (LLMs) for demand sensing (Lauchlan, 2025), offer live tests of this labor-shift claim. Supply chain management (SCM) literature addresses broad GenAI implementation aspects (Wamba

Funding: This work was supported by Ajman University (award no. 2025-IRG-CBA-1).

International Journal of Physical Distribution & Logistics Management
Vol. 56 No. 11, 2026
pp. 241-264
Emerald Publishing Limited
e-ISSN: 1758-664X
p-ISSN: 0960-0035
DOI 10.1108/IJPDLM-08-2025-0444

et al., 2023, 2024; Jackson *et al.*, 2024), operational applications (Li *et al.*, 2024), and performance benchmarks (Dubey *et al.*, 2024), but the socio-technical consequences for SCP professionals remain underexplored.

Labor economics documents the above-mentioned pattern from a different angle. Technology displaces existing tasks while generating new work categories (Acemoglu and Restrepo, 2019), and the net effect on cognitive load depends on whether organizations deliberately redesign jobs around the technology (Parker and Grote, 2022). SCP research on technology-driven system redesign is rich (e.g. Jonsson and Holmström, 2016), but that on role-level consequences is limited. Firms that fail to strategically redesign SCP roles risk overwhelming planners with a dual burden, as the cognitive demands of AI collaboration and validation (Brynjolfsson *et al.*, 2025) accumulate on top of established operational tasks instead of replacing them (Seeber *et al.*, 2020). In other words, planners inherit new responsibilities without shedding old ones, and the resulting role conflict and ambiguity erode the efficiencies that the technology was meant to provide (Tubre and Collins, 2000).

Empirical SCM studies of role-level impacts of GenAI remain scarce, despite the field's recognition of human-centric adoption as a disruptive challenge (see Hendriksen, 2023). SCP presents an ideal empirical setting for this topic, as its core cognitive tasks (e.g. forecasting, scenario analysis, exception handling) map directly onto GenAI capabilities (Brynjolfsson *et al.*, 2025; Huynh, 2024). Firms in the e-commerce, automotive, and consumer packaged goods (CPG) sectors have documented the deployment of GenAI in demand forecasting, inventory optimization, and sales and operations planning (S&OP) processes. Understanding how firms articulate these deployments in SCP role expectations requires data sources that capture both formal requirements and operational rationale.

Job specifications and practitioner interviews fulfill both requirements. Job specifications are strategic artifacts that codify an organization's beliefs about what a role requires, namely its assumptions regarding the necessary skills, responsibilities, and success criteria. They reveal the competencies firms prioritize by acting as physical manifestations of the abstract patterns for a role (Pentland and Feldman, 2005). Triangulating practitioner interviews with documents like job specifications potentially reveals organizational dynamics—in this case, the divergence between formal rules and actual practice—that are often invisible to single-method approaches (Boyer and Swink, 2008).

We combine data from both types of sources in a multiple-case study that compares SCP job specifications across GenAI adopters and non-adopters. The research design separates case classification from case analysis to prevent circular reasoning. Specifically, GenAI adoption status is determined from publicly available strategic evidence (CEO statements, peer-reviewed publications, and third-party analyst reports), while job specifications and interview transcripts are used for the analysis. The analysis explores how documented GenAI adoption corresponds with shifts in required competencies and role orientation. The findings identify two different SCP role archetypes (process guardian vs. architect) and reveal an autonomy–ambiguity paradox that accompanies the adoption of GenAI. We organize these findings into an SCP transformation framework (Figure 2) and a complementary maturity model (Figure 3). These outputs offer an empirically grounded theory of GenAI-driven role transformation for researchers and a diagnostic tool for talent strategy for practitioners.

2. Research method

2.1 Research design and unit of analysis

In its investigation of the impact of GenAI on the SCP profession, this study employed purposeful theoretical sampling. Cases were selected for their potential to illuminate an emergent transformation and provide actionable lessons from the innovation frontier, not for statistical generalizability (Eisenhardt, 1989). Figure 1 maps the four-phase research design: the left-to-right flow follows the methodological logic, wherein classification precedes

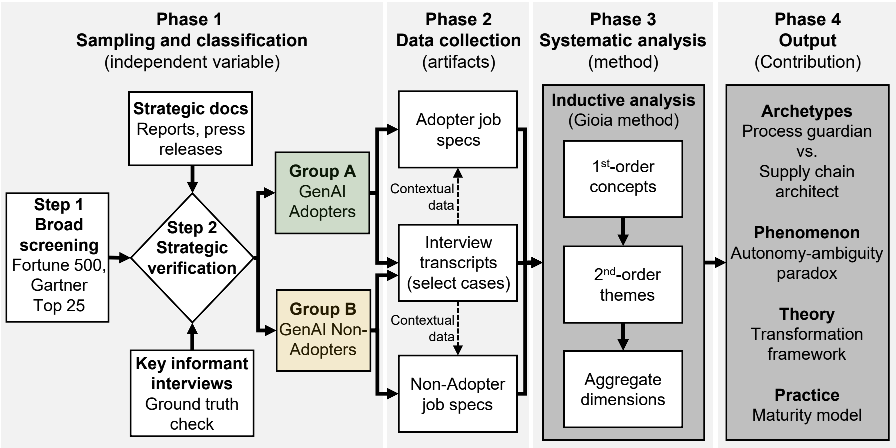

Figure 1. Research design. Source: Authors' own elaboration

collection, collection precedes analysis, and analysis precedes theorization. Phase 1 (sampling and classification) involved two steps. First, we began with a broad screening of firms to generate the candidate pool; then, we performed strategic verification using public documents and key interviews to assign firms to Group A (GenAI adopters) or Group B (non-adopters). This classification occurred prior to analysis of job specifications.

Phase 2 (data collection) collected job specifications from both groups as well as interview transcripts from selected cases, with the latter contextualizing and enriching both adopter and non-adopter analyses (indicated by the dashed arrows) without contaminating the classification logic. In Phase 3 (systematic analysis), we applied the Gioia methodology to the collected data, whereby we progressed from identifying first-order concepts (informant language) to second-order themes (researcher interpretation) and, ultimately, to aggregate dimensions (theoretical constructs). Phase 4 resulted in four contributions: specifically, the analysis identified different role archetypes and an autonomy–ambiguity paradox, which, together with complementary empirical anchors and literature-based insights, helped ground a transformation framework (Figure 2) and an implied practice-oriented maturity model (Figure 3).

This study's primary unit of analysis was formal job specifications for core SCP roles, and the sample comprised positions conventionally classified as tactical-operational. These include demand planners, who generate sales forecasts, and production planners (also known as "master planners" in some manufacturing contexts and "supply planners" in distribution contexts), who translate forecasts into fulfillment schedules and coordinate material requirements given capacity constraints. Also included are integrated roles, such as supply and demand planners or supply chain (SC) analysts, who balance supply strategy (e.g. capacity and inventory) with demand signals (e.g. stock optimization and replenishment parameters). Finally, SC managers overseeing the planning function and S&OP coordinators facilitating cross-functional alignment round out the sample. The sample excluded both purely strategic roles (e.g. vice president of supply chain) and purely executional roles (e.g. warehouse schedulers).

Cross-case analysis matched roles across categories by functional responsibility (e.g. adopter demand planner vs. non-adopter demand planner), industry sector, and hierarchical level. Since GenAI may serve different purposes across contexts (e.g. forecast generation for demand planners, schedule optimization for production planners, scenario synthesis for S&OP coordinators), how GenAI distinctly transforms each role type represents a natural extension of this exploratory work.

2.2 Case sampling and initial screening

The initial pool of potential cases was identified through a systematic review of prominent firms in the SCM and technology sectors. This process involved screening companies from the Fortune 500 list, the annual reports of major publicly traded firms, and the company profiles in leading industry analyst reports, particularly Gartner's SC technology assessments. This screening identified 44 candidate firms in the e-commerce, automotive, CPG, retail, and technology sectors.

Two criteria governed selection from this initial pool. The first was data accessibility, whereby firms needed to have publicly available SCP job specifications on official career portals or major job boards during the window of data collection. The second was strategic positioning, in that firms had to have the potential to be classified as either technology innovators or established industry leaders, based on preliminary evidence from annual reports and analyst coverage. This criterion ensured relevance to the research question and comparability, as identifying transformation at the innovation frontier requires sampling both firms positioned as early adopters and established industry leaders to provide the comparative baseline.

The selection process did not rely on the content of the job specification. A firm's inclusion in the study depended entirely on its position as laid out in external strategic documents, not on what its job postings said about AI or technology. This separation preserved the validity of the subsequent analysis.

2.3 Classification protocol

The classification of firms as GenAI adopters (Group A) or non-adopters (Group B) followed a protocol designed to prevent circular reasoning. Classification data had to remain strictly separate from analysis data, a separation we termed the "methodological firewall." Accordingly, strategic evidence informed the status of GenAI adoption (the independent variable), while job specifications informed the identification of SCP role changes under GenAI adoption (the dependent variable).

For classification into Group A, strategic evidence had to document that generative AI technologies, namely, systems capable of producing novel outputs from learned distributions (including LLMs, transformer-based forecasting architectures, RAG systems, and agentic AI frameworks built on generative foundations), were deployed or actively built in the SCP function. Conventional machine learning (ML) models performing classification, regression, or optimization without generative output capability did not satisfy this criterion on their own. Furthermore, for agentic AI frameworks in the evidence base, classification as GenAI deployment required confirmation that the underlying reasoning engine was a generative foundation model; conventional rule-based or optimization systems that coordinate multi-step execution did not satisfy this criterion.

Three categories of strategic evidence were considered. First, CEO and executive statements from annual reports, press releases, and earnings calls (quarterly investor briefings in which executives discuss company performance) represented evidence with the highest authority. These were followed by technical publications (peer-reviewed papers, Amazon Science articles, corporate engineering blogs) and vendor case studies (published implementations by technology providers, such as Kinaxis documenting client deployments) that confirmed specific GenAI deployments. Finally, industry-analyst validations from sources such as Gartner provided independent third-party confirmation. Annex A1 catalogs the specific sources used to establish the GenAI adoption status for each firm in this category.

The strategic evidence was structured into two tiers. Tier 1 comprised explicit GenAI mentions with SCP-related functional specificity, i.e. instances where executives or technical publications named GenAI, LLMs, or architecturally equivalent technologies (transformer-based forecasting models) in demand planning, production planning, inventory management, or S&OP contexts. Amazon's deployment of Chronos (a T5-based generative forecasting model published in Amazon Science) and Colgate-Palmolive's C-suite confirmation of GenAI and ML for several years in their demand planning cycles exemplify this tier. Tier 2 consisted of mentions of GenAI in adjacent operational domains (e.g. manufacturing intelligence, robotics coordination, product development) for which the capabilities extend to SCP functions. Tesla explicitly hiring data algorithm engineers with LLM and multimodal experience for SCM roles (which were not part of the SCP roles analyzed in this study) exemplifies this tier.

Annex A2 synthesizes the strategic evidence strength assessment for each adopter firm and documents the deployment architecture distinctions. The classification data revealed meaningful variations in how firms deploy GenAI. For instance, Amazon and Colgate-Palmolive operate GenAI-native architectures, in which generative models are embedded in core planning algorithms, such as transformer-based forecasting and LLM-driven demand sensing. The Warehouse Group operates GenAI-augmented architectures, wherein the Budget Bot functions as a decision-support layer atop ML-based forecasting operations. Tesla's evidence base confirms active LLM capability-building in engineering roles with SCP functions, but it does not explicitly detail operational deployment. Given these variations, Annex A2 delineates three deployment architectures within Group A: GenAI-native (Amazon, Colgate-Palmolive), GenAI-augmented (The Warehouse Group), and build-phase (Tesla).

Non-adopter classification followed the inverse logic [2]. A firm qualified as a non-adopter when a thorough search of annual reports, analyst coverage, and industry databases yielded no credible evidence of GenAI deployment in SCP functions. The burden of proof fell on demonstrating this absence. We validated this classification through direct outreach. To this end, LinkedIn messages were sent to one or more randomly selected SCP practitioners at each non-adopter firm for inquiring whether GenAI tools supported their planning tasks. Responses from at least one representative per non-adopter confirmed that either no GenAI tools were in use or that experimentation with these tools was limited to some individuals and without organizational integration or support. Table 1 presents the resulting matched pairs by industry sector and scale.

Table 1. Representative matched-pair case profiles for the cross-case analysis

SCP role category	Role description	GenAI adopter (Group A)	GenAI non-adopter (Group B)	Matched industry
Demand planning	Generates sales and demand forecasts	Amazon[a]	Pattern	E-Commerce
		Tesla	Bugatti Rimac	Automotive
		Warehouse Group	Gottardo	Retail
Production planning	Translates forecasts into production schedules and manages material supply	Tesla[a]	GKN Automotive	Automotive
		Colgate-Palmolive	EssilorLuxottica	CPG
Integrated supply and demand planning	End-to-end balancing of demand signals with inventory and supply strategy	Amazon[a]	Shopee	E-Commerce
		Colgate-Palmolive	ASICS	CPG
S&OP	Facilitates strategic alignment across commercial, financial, and operational functions	Amazon[a]	Seagate	Technology

Note(s): Consumer Packaged Goods (CPG). Classification of GenAI adopter/non-adopter was based on strategic evidence (Annexes A1–A2), [a]Cases selected for the longitudinal analysis of supply chain planning role evolution

Source(s): Authors' own elaboration

2.4 Data collection

Data was collected in two phases. In Phase 1, the primary data consisted of full-text job specifications for core SCP roles, sourced from official career portals and major job boards (e.g. LinkedIn, Indeed). For the cross-sectional analysis, specifications posted within three months of the data collection cutoff were collected for both Group A and Group B firms. For the longitudinal analysis, corresponding job descriptions from approximately eight months prior were collected for selected Group A firms (Amazon, Tesla, Colgate-Palmolive), to track within-firm role evolution. Annex A4 provides metadata for each collected specification.

In Phase 2, four in-depth interviews (45–60 min each) with senior SCP leaders provided triangulation for the job specification analysis. Job postings demonstrated what competencies firms seek, whereas interviews revealed why those competencies matter and how planners deploy them. The interviewees from Group A were a senior manager at Amazon, a supply chain manager at Tesla, and a director at Colgate-Palmolive; the interviewee from Group B was a production planning manager at GKN Automotive. The interview protocol (Annex A9) addressed the evolution of planning tools, the impact of automation, skill requirements, and organizational dynamics surrounding autonomy and ambiguity. The questions were designed to elicit operational aspects without steering interviewees to confirm preexisting classifications.

The interview data served two functions. First, for adopters, they provided ground-truth checks for strategic evidence. For example, the Amazon interview confirmed the existence of an internal GenAI tool that is not documented in public documents, while the Tesla interview revealed the use of advanced AI for predictive disruption detection. Second, for analytical enrichment, all four interviews revealed mechanisms that were not visible in job specifications. For instance, the Amazon interview quantified productivity gains that no job specification mentioned. Similarly, the Colgate-Palmolive interview described a shift in their S&OP process from opinion-based negotiation to data-driven debate, which affected planners' identities and was not reflected in the corresponding job specifications. Annex A3 synthesizes the interview data by thematic dimension.

2.5 Analytical approach

The analysis entailed a systematic integration of the two data streams. The job specifications were subjected to qualitative content analysis, guided by an initial coding scheme from the literature (role conflict, technical skills, strategic orientation) and by emergent themes from the data (Strauss and Corbin, 1998). This analysis enabled a systematic comparison both between groups and over time (longitudinally). Annexes A5 and A6 present detailed cross-case and longitudinal comparisons. The thematic structure in Table 2 was derived using the Gioia *et al.* (2013) methodology. First-order codes remained faithful to terms used in the source materials (e.g. "proactive risk identification," "systems thinking and innovation"). These codes were then clustered into second-order interpretive themes (e.g. "proactive systems orchestration," "advanced and programmatic tools"). Thereafter, we consolidated those themes into aggregate analytical dimensions ("core mission and philosophy," "essential tools and platforms," "competencies and mindset"). In this analysis, we did not consider the mention of traditional tools (e.g. SQL, Excel, ERP) in adopter job specifications to be contradictory to a GenAI classification, as these tools can serve different purposes across competency clusters. For instance, in non-adopter contexts, SQL extracts data for manual analysis, whereas in adopter contexts, SQL feeds the data pipelines that support AI-generated insights; adopter specifications pair SQL with ambiguity tolerance, systems thinking, and collaborative influence, while non-adopter specifications pair SQL with transactional accuracy, process discipline, and detail orientation. This pattern is evident in Table 2, which shows that SQL proficiency appears in job specifications for Amazon, Colgate-Palmolive, and EssilorLuxottica (a non-adopter), but the surrounding competency clusters differ. Table 2 synthesizes the competency analysis, where each checkmark can be traced to specific language in the source specifications documented in Annex A4.

The interview transcripts were subjected to parallel thematic analysis, so that the two datasets could be synthesized into a coherent narrative. Job specifications revealed which skills and responsibilities are changing, and interviews explained why and how. The findings presented in Section 3 integrate both streams.

3. Evolution of the SCP role

Two distinct role configurations—distinguished by mission, tools, and competencies—emerge from the analysis: one in non-adopter firms, the other in adopter firms. Table 2 synthesizes the divergence between these archetypes. The following sections explore these patterns.

3.1 Divergence in core mission and philosophy

The fundamental difference between the archetypes lies in their guiding philosophies (Table 2). The non-adopter archetype is oriented toward exception-driven execution, or responding to arising deviations within established parameters. The adopter archetype is oriented toward orchestrating anticipatory systems, or designing planning processes that sense and adapt to emerging conditions. This distinction concerns role orientation, regardless of the uncertainty-regulation strategy employed. Reactive postures can be strategically appropriate for certain types of uncertainty (Sengupta *et al.*, 2025); the limitation of the non-adopter archetype is not reactivity *per se*, but rather the absence of system-design authority.

Three attributes define the mission of the non-adopter archetype: process discipline ensures adherence to established procedures; transactional accuracy guarantees data integrity; and exception management addresses deviations as they arise. Success is measured by adherence to predefined targets, such as on-time in-full delivery (Christopher, 2016). The GKN Automotive interviewee articulated this philosophy: “The only way to control a process is to measure its outputs and find the deviations . . . You’ll also need to understand their root causes and then correct them . . . All of our data analysis is in service of that loop.”

The mission of the adopter archetype centers on orchestrating outcomes across multiple types of uncertainty. Systems thinking enables planners to identify interdependencies across the supply network, and end-to-end strategic ownership extends accountability beyond functional boundaries. Risk identification shifts attention from reacting to exceptions toward anticipating them. The Amazon manager described the objective as “turning chaos into predictability.” The Tesla interviewee articulated the shift from execution to design: “I don’t plan optimally. I engineer the flow . . . I dismantle constraints to their atomic parts and then reconstruct the supply chain path to eliminate those constraints.” In adopter firms, the value of the planner role is derived from their ability to design and improve the system itself.

This mission contributes to enhanced SC agility, represented by an alertness in sensing market shifts, a decisiveness in formulated responses, and a swiftness in timely implementing them (Gligor *et al.*, 2013). The adopter archetype manages this entire sense-and-respond cycle by engineering a system to detect and respond to emerging conditions before they crystallize into exceptions. This transition from a reactive to an anticipatory work posture corresponds with documented patterns in adjacent AI-enabled supply chain functions (Ciceri *et al.*, 2026); in those contexts, human analysts shift from exception-response to proactive risk treatment as AI handles the routine pattern recognition.

This expanded mandate, however, contains a paradox. GenAI increases planner autonomy. The Amazon interviewee described planners making “million-dollar inventory decisions on their own” that previously required senior approval. However, the same technology increases ambiguity, and the Amazon interviewee quantified the cognitive shift: “We’ve moved from a world of known unknowns to a world of unknown unknowns . . . Before, maybe you’re analyzing three scenarios for a big event, but now, AI can generate three hundred scenarios with different probability weights.” “Probability weights” are AI-generated confidence scores that reflect the model's internal assessment of its own scenario outputs. Historical frequency distributions produce no such scores, which is precisely what makes scenario selection a judgment problem. Calculating probabilities is the machine's task; the planner's ambiguity lies in determining which AI-generated scenario warrants action, even when the underlying uncertainty may be genuinely unknowable.

In short, planners gain decision authority as the decision space becomes harder to define. We term this the *autonomy–ambiguity paradox*, which expands paradox theory's emphasis on tensions requiring management rather than resolution (Smith and Lewis, 2011) and extends the automation–augmentation paradox in AI management (Raisch and Krakowski, 2021) to the specific cognitive demands of GenAI-enabled planning. The paradox is structural and so does not resolve as planners gain experience. Earlier waves of automation reduced both autonomy and ambiguity through process standardization, but GenAI produces the opposite effect by expanding the decision space while increasing planner authority within it. Manual analysis constrained planners to a handful of scenarios, which artificially narrowed the visible decision space. GenAI removes that constraint and exposes the range of plausible futures that manual methods failed to show. As these systems mature, they reveal finer contingencies, so that judgment challenges deepen with the technology. This rising demand for judgment under expanding ambiguity defines the adopter archetype.

3.2 Divergence in essential tools and platforms

Mission divergence corresponds with tool divergence. Non-adopter firms build their planning functions on transactional enterprise resource planning (ERP) and Office suites (Table 2), while adopter firms provide their planners with advanced programmatic toolkits.

Planners in non-adopter firms operate within a transaction-focused environment, even where ERP platforms carry embedded analytics and AI capabilities. Deployment intent explains the difference. Non-adopter firms use ERP primarily as a records system for order processing, inventory transactions, and production scheduling. The GKN Automotive interviewee described this reliance: "Without a doubt, the single piece of software that would make a planner's job [at GKN] impossible [without it] is our SAP ERP system . . . The second . . . and I know this will sound old-fashioned . . . Excel." Job specification data substantiate this pattern, with repeated emphasis on foundational ERP mastery and advanced Excel proficiency in requirements across non-adopter cases. This transaction-focused orientation characterizes an early stage of SC digitalization, where the primary objective is process standardization and data integrity (Trkman *et al.*, 2010).

Some non-adopter specifications also list newer tools. For instance, EssilorLuxottica requires Python, and Seagate specifies advanced planning systems (APSs). This seeming overlap with adopters requires interpretation. Simply put, tool presence differs from tool function. For example, EssilorLuxottica's Python requirement appears alongside demands for detail orientation, transactional accuracy, and tactical supplier relationship management (Table 2); here, the tool supports efficiency optimization in an execution-focused role. Amazon's Python requirement appears alongside demands for ambiguity tolerance, systems thinking, and collaborative influence, whereby the tool serves an architectural function within a strategic orchestration role. The distinguishing factor is the competency cluster surrounding the tool, not the tool itself. Empirical evidence from human-in-the-loop AI systems reinforces this point, as explainability infrastructure drives adoption outcomes as strongly as predictive accuracy does (Sridhar *et al.*, 2026).

Job specifications from adopter firms downplay ERP proficiency in favor of programmatic analytical skills. The Amazon interviewee stated, "Python is the new Excel, straight up. If someone can't write Python scripts or at least understand them to automate workflows or build quick analyses, they're gonna struggle . . ." Python enables the pipeline architecture that feeds data into AI models, automates queries, and extracts outputs from APSs; proficiency matters because of what it makes possible downstream. This shift moves beyond data manipulation

Table 2. Detailed synthesis of role competencies from job specifications

Dimension	Sub-dimension	Attribute	GenAI adopter Amazon	Tesla	Colgate-Palmolive	Warehouse Group	GenAI non-adopter GKN automotive	EssilorLuxottica	Bugatti Rimac	Pattern	Seagate	Asics	Gottardo	Shopee
Core mission and philosophy	Systems orchestration	Systems thinking and innovation	X	X	X	X								
		End-to-end strategic ownership	X	X	X						X			
		Risk identification and mitigation	X	X	X	X								
	Exception-driven execution	Process discipline and adherence to plan					X	X	X	X		X	X	X
		Transactional accuracy					X		X	X	X	X	X	X
		Exception management and firefighting					X	X		X	X	X		X

(*continued*)

Table 2. Continued

			GenAI adopter				GenAI non-adopter							
Dimension	Sub-dimension	Attribute	Amazon	Tesla	Colgate-Palmolive	Warehouse Group	GKN automotive	EssilorLuxottica	Bugatti Rimac	Pattern	Seagate	Asics	Gottardo	Shopee
Essential tools and platforms	Advanced and programmatic tools	SQL proficiency	X		X						X			
		Python/ programmatic scripting	X					X						
		Business intelligence and visualization (Tableau, Power BI)	X	X		X		X		X	X			
		Advanced planning systems (APS)			X	X					X			
		Explicit mention of AI/ML/LLMs			X									
	Traditional ERP and Office suites	Advanced Excel (modeling, macros)	X	X	X	X	X	X	X	X	X	X	X	X
		Foundational ERP mastery			X		X	X	X	X	X	X	X	X

(*continued*)

Table 2. Continued

Dimension	Sub-dimension	Attribute	GenAI adopter				GenAI non-adopter							
			Amazon	Tesla	Colgate-Palmolive	Warehouse Group	GKN automotive	EssilorLuxottica	Bugatti Rimac	Pattern	Seagate	Asics	Gottardo	Shopee
Competencies and mindset	Analytical and engineering mindset	First-principles/ engineering logic		X										
		Ambiguity tolerance/ navigating uncertainty	X	X	X									
		Collaborative influence and stakeholder management	X	X	X	X					X			X
		Project/program management	X	X	X	X					X			
	Procedural and executional mindset	Detail orientation and accuracy			X	X	X	X	X	X	X	X	X	X
		Tactical communication and coordination					X	X	X	X	X	X	X	X
		Tactical supplier relationship management		X			X	X	X		X	X	X	X
		Experience in high-volume environments		X	X	X	X	X				X		X

Note(s): An X indicates that the attribute was explicitly mentioned or strongly implied as a key requirement in the analyzed job specifications for the corresponding case. Annex A4 provides the source specifications; Annex A5 presents the detailed cross-case comparison underlying this synthesis

Source(s): Authors' own elaboration

with existing tools to programmatic creation of new analytical processes. Deploying advanced AI requires organizational capabilities, such as managing unstructured data and developing predictive models (Richey *et al.*, 2023).

The interviews also identified GenAI tools not listed in job postings. The Amazon interviewee described an internal GenAI system accessed by planners through natural language prompts: "She literally just prompts our [Gen]AI tool with something like "analyze demand anomaly for [certain item number], consider social media trends ... competitive landscape ... and create demand scenarios for the next eight weeks with confidence intervals."" Analyses that previously required 6–8 h of manual data extraction and scenario construction now take 5 min. The Colgate-Palmolive interviewee described similar automation, wherein the planner receives system-generated forecasts accompanied by a "[Gen]AI[-powered] narrative" explaining the predicted promotional lift based on historical performance, competitor activity, and marketing support. This gap between documented practice and codified hiring expectations represents a documentation lag, which occurs when formal documentation evolves more slowly than operational practice (Atalay *et al.*, 2020), as HR departments update role specifications through periodic reviews (Kulik and Perry, 2023). To wit, Amazon and Colgate-Palmolive deploy GenAI tools that have reshaped planner workflows, yet their postings continue to emphasize SQL and Excel alongside emerging requirements such as Python and AI literacy.

Colgate-Palmolive's platform evolution illustrates this shift from transaction-focused ERP toward AI-integrated planning. The interviewee explained: "The first major shift ... was our strategic decision to partner with Kinaxis for our advanced planning solution. This was a monumental change for us ... we were buying into concurrent planning." The firm's specifications now explicitly list Kinaxis as a requirement, along with experience with AI and digital tools. The platform choice aligns with evidence that integrated planning systems with concurrent capabilities improved firm performance in SC disruptions (Swink *et al.*, 2025).

3.3 Divergence in competencies and mindset

Mission and tool divergence correlate to competency divergence. On this point, Table 2 reveals the distinction between procedural and analytical engineering mindsets.

Non-adopter job specifications frequently require attention to detail, accuracy, and tactical communication and coordination skills. SCM literature has historically defined effective planners using these competencies, with an emphasis on professionals skilled in data handling, communication, and process execution within established ERP frameworks (Flöthmann *et al.*, 2018). The GKN interviewee reinforced this profile: "So, no, I don't think I'll be hiring for a fundamentally different kind of planner. I'll still need planners who are disciplined and have a sharp eye for detail."

Data from adopter job specifications and interviews reflect a substantive evolution beyond procedural expertise. Executional skills remain important, but the defining competencies center on an analytical engineering mindset. Three attributes distinguish this profile.

The first defining attribute, first-principles reasoning, requires planners to strip problems back to foundational constraints (cycle time, tooling capacity, logistics parameters) and to construct solutions from that base. GenAI outputs hold no inherent awareness of such constraints, so validation falls to the planner. The Tesla interviewee framed the attribute as a hiring necessity, with the ideal candidate being a "hybrid ... part engineer, part planner, part analyst ... but [with] the mindset of an owner and problem-solver." The interviewee elaborated: "We need people who think in first principles [who understand cycle time, geometry, tooling capacity, logistics], not just Excel macros ... If you can scope a fixture improvement, set up a logistics workaround, code a quick SQL query to validate delivery data ... then run in a twenty-four-hour sprint ... you're the kind of person I hire." This reflects increasing complexity in digital SC, for which planners need system-level understanding to manage digital twins and interconnected data flows (Ivanov *et al.*, 2019).

The second defining attribute is ambiguity tolerance. This competency addresses the scenario-selection paradox mentioned in Section 3.1, wherein planners who managed three scenarios must now evaluate hundreds for which the cognitive challenge is not computational but judgmental (i.e. selecting among AI-generated alternatives when historical patterns offer no reliable guide). Greater data visibility does not necessarily reduce complexity (Tan *et al.*, 2015). This challenge may extend beyond scenario volume to situations where no precomputed scenario applies, and planners must thus improvise responses to unprecedented conditions. In such dynamic SC environments, managers navigate complexity through continuous sensemaking by interpreting data, adapting routines, and making judgments in ever-changing contexts (Wieland, 2021).

The third defining attribute is collaborative influence. As planning becomes more automated, planner value shifts from calculation to facilitation. The Colgate-Palmolive interviewee quantified this: “A decade ago, the balance might have been sixty to forty in favor of technical skills. Today, it’s completely flipped. I would say it’s at least seventy to thirty in favor of the ability to collaborate and understand the business.” For Colgate-Palmolive, technical proficiency has become a qualification threshold, and what differentiates candidates lies beyond it: “Having Kinaxis on your resume will get you through the door, but it won’t get you the job. We have found that we can teach a smart analytically-minded person the specific clicks and navigation of our tools. But it is far more difficult to teach someone business aspects or how to build a trusted relationship with sales.” Accordingly, job interview criteria have shifted: “I’m much more interested in how a candidate resolved a disagreement on a forecast than I am in how they built it.” Here, the planner evolves into a central S&OP actor responsible for orchestrating cross-functional consensus and influencing strategic decisions (Jonsson and Holmström, 2016).

It could be argued that these competency patterns reflect a general analytical sophistication that is broader than GenAI-specific demands. Our rebuttal involves three forms of evidence from our analysis distinguishing GenAI effects from broader quantitative intensity and pointing to a GenAI-specific explanation. First, the documented competency attributes of ambiguity tolerance, first-principles reasoning, and scenario selection address cognitive challenges unique to generative systems. Evaluating machine-generated scenarios, validating probabilistic recommendations, and selecting among AI-generated alternatives differ from traditional quantitative demands. Second, interview data identify the specific mechanisms driving these competency requirements. Amazon’s internal tool and Colgate-Palmolive’s automated narrative generation are transformer-based and LLM-driven capabilities that reshape planners’ cognitive tasks, while Tesla’s job specifications contain explicit LLM and GenAI requirements for engineering roles in SCP functions, which reflect active investment in equivalent capabilities. Third, the timing of strategic evidence (Annex A6) supports a causal direction. Amazon added AI literacy requirements after the 2023 deployment of Chronos (a transformer-based forecasting model) and COSMO (an LLM for demand signal enrichment), which supports a GenAI-driven explanation over a general sophistication alternative.

4. Discussion

We synthesize the identified patterns of SCP archetypes (seen in divergent missions, tools, and competencies) into a transformation framework that models how GenAI adoption progresses toward either augmentation or overwhelm, which we operationalize as a maturity model with diagnostic stages and transition requirements. Managerial implications and a research agenda are discussed in subsequent sections.

4.1 Transformation framework

Figure 2 models the transformation process across three zones. The disruption point is the deployment of GenAI, which creates the autonomy–ambiguity paradox, whereby planners gain

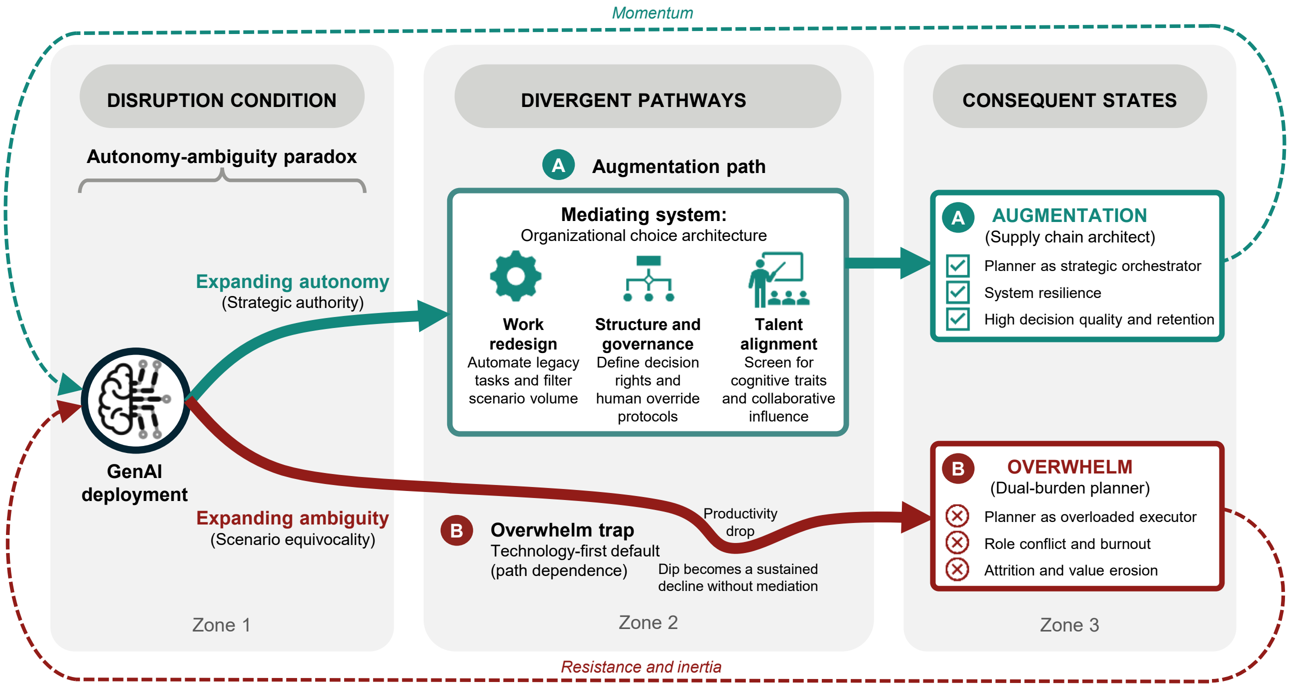

Figure 2. GenAI-driven supply chain planning transformation framework. Path A (teal) represents strategic intervention through organizational choice architecture; Path B (red) represents the technology-first default that bypasses mediating factors. Solid arrows indicate primary causal paths. Dashed feedback loops indicate path dependence, with augmentation success generating momentum for further transformation and overwhelm generating resistance that inhibits future adoption. The augmentation path reflects observed case dynamics, while the overwhelm path represents a theoretically derived trajectory that is consistent with role conflict literature and documented patterns in technology implementation without deliberate organizational redesign. Source: Authors' own elaboration

strategic authority over decisions previously reserved for senior management while facing ambiguous situations in which no historical pattern applies and no meaningful probability can be assigned. Paradoxes require management, not resolution (Smith and Lewis, 2011), and Paths A and B diverge based on whether organizations manage this paradox. Generative AI shifts human labor from content production to verification and curation, and agentic AI reconfigures decision authority by delegating multi-step task execution (Chalmers *et al.*, 2026). In GenAI-native SCP firms, the two mechanisms operate simultaneously, with the former constituting the ambiguity arm of the paradox and the latter constituting the autonomy arm. The same deployment event that grants authority creates the cognitive condition that makes its exercise demanding.

Path A represents strategic intervention through organizational choice architecture (Thaler *et al.*, 2013). We propose three intervention categories that collectively form a system that mediates the effect of GenAI deployment on organizational outcomes. The first category is SCP work redesign. Technology alone does not improve job performance (Parker and Grote, 2022); SCP must be redesigned to offload routine tasks to automation while filtering AI outputs into manageable forms. The second category is governance. Redesigned workflows create new decision points that require clear authority structures (Aghion and Tirole, 1997), as planners need clarity on who accepts AI recommendations, who can override them, and what triggers escalation. The third category is talent strategy. New workflows and governance structures demand different competencies. To this end, organizations should prioritize learning potential (Xiang *et al.*, 2024) by treating the competencies identified in Section 3.3 (ambiguity tolerance, first-principles reasoning, and collaborative influence) as screening criteria and AI literacy development as a post-hiring imperative.

Implementing this system is not frictionless, and we acknowledge that navigating new collaboration routines may trigger an initial productivity dip (Brynjolfsson *et al.*, 2021). Nevertheless, the Amazon interviewee confirmed that this dip is transient, provided that planners receive adequate support (concerning the interventions' purpose). Once stabilized, planners can function as strategic orchestrators, with higher decision quality and improved retention. Phased progression is also consistent with a six-year longitudinal study of AI deployment across 30+ projects at Robert Bosch (Canciglieri *et al.*, 2026); in that study, initial adoption did not yield immediate improvements, and premature intervention in early phases proved detrimental. In short, organizational maturation through AI deployment follows non-linear trajectories.

Path B is the technology-first default that we call the "overwhelm trap." In this scenario, an organization deploys GenAI without the mediating system; the inevitable productivity dip still occurs, yet management may misinterpret it as a failure of implementation or as planner resistance. Change resistance in AI adoption is structurally induced by gaps in infrastructure and skill readiness (Mandal *et al.*, 2026), and the organizations that generate resistance create it through the sequence they choose. At this point, the SCP environment deteriorates. Scenario volume grows without filtering, decision rights stay unclear, and competencies lag behind task demands. The initial dip hardens into a sustained decline, and the result is a sense of overwhelm as planners bear the double burden of established responsibilities and new AI demands. This burden may lead to role conflict, burnout, and attrition, which erode both the talent base and the strategic value that GenAI was intended to unlock.

These two trajectories create distinct professional identities. On Path B, planners remain locked in the same execution-focused configuration observed among non-adopters, we call these process guardians. Their value depends on procedural mastery and is defined by transaction accuracy, data integrity, and exception management within established parameters. Path A transforms planners into what we call SC architects. Here, value flows from system design, defined by AI configuration, uncertainty navigation, and cross-functional integration. GenAI expands SCP capability across all four quadrants of uncertainty navigation (Sengupta *et al.*, 2025): predictive optimization for known-known variation, proactive sensing for known-unknown patterns, reactive protocols for unknown-known events, and adaptive judgment for unknown-unknown emergence.

Path dependence governs both trajectories (Sydow *et al*., 2009). Successful augmentation builds confidence that accelerates subsequent transformation, whereas overwhelm breeds defensive behaviors that impede future adoption. The feedback loops in Figure 2 capture these self-reinforcing dynamics.

Tesla's classification as a build-phase adopter represents a theoretically distinct organizational state that this two-path framework does not explicitly anticipate. Strategic commitment to GenAI capability investment precedes operational deployment, and the role pressures in confirmed adopters (as documented by this study) are downstream of that deployment. Examining Tesla alongside GenAI-native firms and non-adopters enabled separating the antecedents of adoption from its consequences, such that strategic investment decisions, talent restructuring, and engineering capability-building could be studied before the paradoxes they eventually generate arise. Which trajectory a transitional firm ultimately follows remains an open empirical question.

4.2 Maturity model

Figure 3 provides a maturity model for determining organizational position and identifying pathways forward. The model complements AI capability frameworks that map technology types to operational contexts (e.g. Richey *et al*., 2023) by focusing on how human roles co-evolve with such capabilities. The five levels in each of the three dimensions enable practitioners to diagnose their current position and identify the transition paradox impeding progress.

The first dimension, role orientation, concerns how the planning function defines its core mission. Level 1 is focused on executing externally generated plans, whereas level 5 is focused on managing AI agents and designing adaptive systems. The progression moves from execution through optimization, advisory, and ownership to orchestration.

The second dimension, the tool ecosystem, defines the technological infrastructure within which planners operate. Level 1 relies on transaction-focused ERP and manual spreadsheets; at the other end of the spectrum, level 5 features real-time proprietary models (e.g. custom forecasting algorithms trained on firm-specific data) and digital twins (virtual replicas of supply network nodes that simulate disruption scenarios before they occur). The progression moves additively from ERP-centric systems through business intelligence (BI)-enhanced dashboards, programmatic scripting, and APS-integrated platforms to AI-native architectures. For example, a level 4 organization still runs ERP for transactional integrity, but it layers APS platforms and AI tools on top of that. What shifts for earlier tools is the emphasis. ERP evolves from a central planning system into a data source that feeds AI-driven planning. This layering explains the documentation lag identified in Section 3.2 (i.e. formal job specifications trailing practice), as GenAI adopters may still emphasize ERP proficiency even when they are operationally dependent on GenAI tools.

The third dimension, competency profile, represents the valued skills at each level. Level 1 prioritizes procedural mastery, defined by discipline, accuracy, and attention to detail. Level 5 requires collaborative intelligence, which involves discerning when to trust AI recommendations, when to override them, and how to exercise judgment when algorithmic outputs conflict with contextual knowledge (Wilson and Daugherty, 2018). Here too, the progression is additive; level 5 planners still need procedural discipline, as transaction errors can cascade through AI systems just as they do through manual processes. Here, the shift is one of relative weight. Procedural mastery qualifies candidates for consideration, while collaborative intelligence determines who advances.

The progression to each subsequent level is anchored in the study's empirical data. Non-adopter job specifications (from GKN, EssilorLuxottica, and Pattern) cluster at levels 1–2, as they emphasize process discipline and ERP mastery (Table 2). Adopter specifications (from Amazon, Tesla, and Colgate-Palmolive) cluster at levels 4–5, as they prioritize ambiguity tolerance and AI literacy. Annex A7 provides detailed specifications, diagnostic indicators, and verbatim empirical anchors.

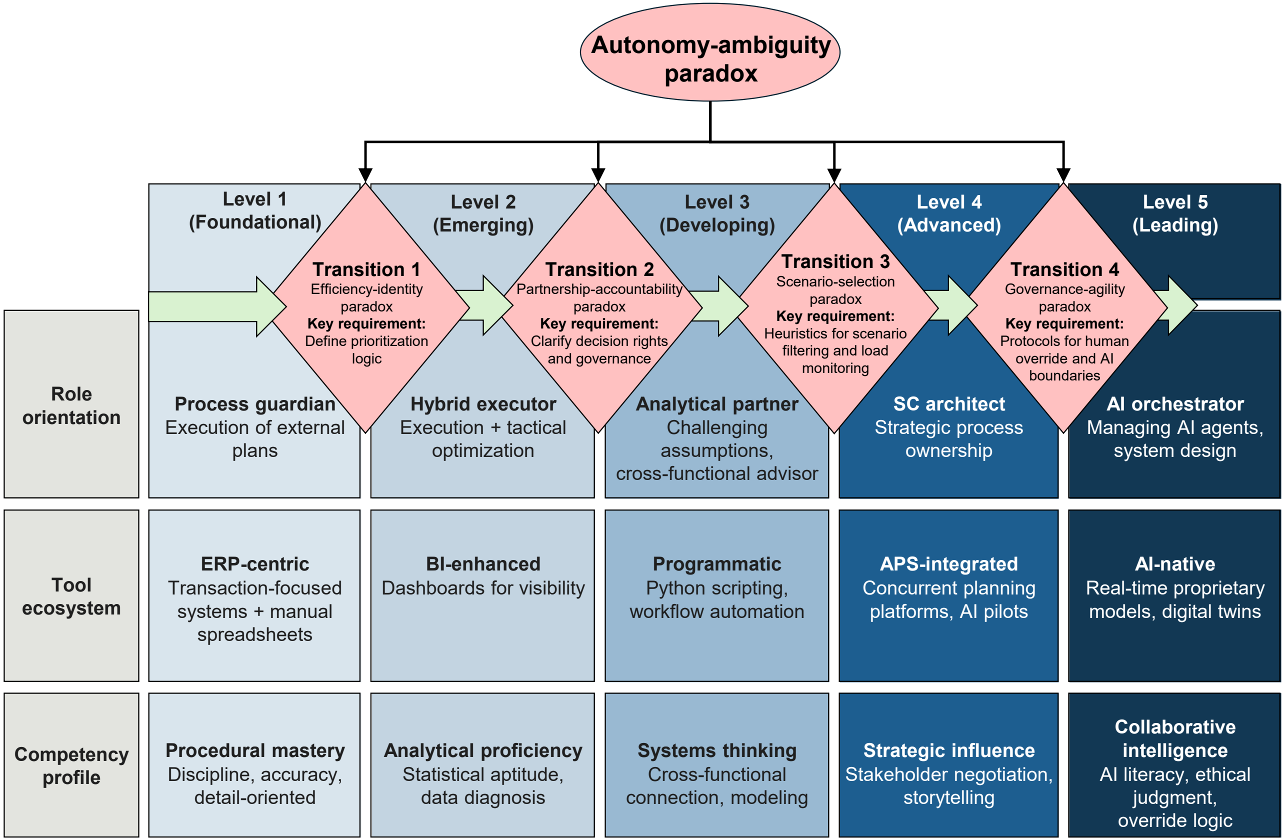

Figure 3. SCP function maturity model. Levels 1–2 characterize non-adopter configurations, and Levels 4–5 characterize adopter configurations. Progression is additive—i.e. Higher levels incorporate new tools and competencies while retaining foundational capabilities. Diamond elements identify transition-specific manifestations of the autonomy–ambiguity paradox with the associated resolution requirements. Detailed specifications and empirical anchors appear in Annexes A7–A8. Source: Authors' own elaboration

Progression through the model is not linear, and in the transitions between levels (diamond elements in Figure 3), specific forms of the autonomy–ambiguity paradox arise. Annex A8 details each transition's manifestation, empirical anchors, and resolution mechanisms.

The move from level 1 to level 2 triggers an efficiency–identity paradox, in which automation threatens the execution identity that defines process guardians. Consequently, planners resist improvements that eliminate their core tasks. Management must address this by defining a prioritization logic concerned with freeing capacity (Levinthal and Wu, 2025).

The shift to level 3 entails a partnership–accountability paradox, in which influence expands as planners assume advisory functions, yet formal authority lags behind. As a result, planners bear responsibility for recommendations they cannot enforce. Resolving this tension requires explicit clarification of decision rights and governance structures (Puranam, 2021).

Advancing to level 4 introduces a scenario–selection paradox (Section 3.1). The sheer wealth of AI options increases cognitive burden, as planners must select from potentially hundreds of choices while matching suitable responses to various uncertainty types (Sengupta *et al*., 2025). Resolving this burden requires filtering heuristics (Bettis, 2017) and monitoring whether scenario volume exceeds planner processing capacity (Levinthal and Wu, 2025), since the bottleneck is attentional and additional analytical capacity leaves it untouched. The goal is to establish a continuous sensemaking practice that is necessary for navigating dynamic contexts (Wieland, 2021).

Finally, the leading edge (transition from level 4 to level 5) presents a governance–agility paradox, in which oversight requirements constrain the very responsiveness that justifies AI adoption. This tension demands protocols that specify human-override authority and AI operational boundaries (Shrestha *et al*., 2019).

4.3 Managerial implications

Generic advice to redesign roles or to upskill talent for effective AI adoption provides insufficient guidance. In contrast, the maturity model specifies which interventions matter at each stage. Practitioners should assess their current level, identify the specific transition paradox impeding them, and design interventions that target that exact obstacle.

Organizational design requirements change as firms advance. At transitions $1 \rightarrow 2$ and $2 \rightarrow 3$, work redesign focuses on automating specific legacy tasks, such as manual data extraction and routine exception routing, while role redefinition involves formalizing efficiency improvement responsibilities and establishing initial governance boundaries so that planners understand their new mandate (Parker and Grote, 2022). Sequencing matters. Infrastructure readiness must precede skill development, and skill development must precede behavioral adoption (Mandal *et al*., 2026). Organizations that invest in AI platforms before addressing these foundations encounter the structural resistance delineated by the overwhelm path.

At transition $3 \rightarrow 4$, the complexity increases. Here, work redesign entails implementing scenario-filtering protocols calibrated to the type of uncertainty. Predictable demand variation (known-known) informs automated optimization, while ambiguous patterns (known-unknown) are referred for human review with AI-generated context, and genuine disruptions (unknown-known, unknown-unknown) are escalated immediately. This approach is illustrated by Amazon's internal system, where hierarchical scenario tiers with automated prescreening only route scenarios that exceed deviation thresholds for human review. In addition, at this transitional stage, role redefinition requires establishing decision rights for AI recommendation acceptance and override. Planners need clarity on when to follow system recommendations, when to deviate from them, and which types of uncertainty require immediate human judgment.

As organizations reach transition $4 \rightarrow 5$, governance becomes the primary concern. Human–AI collaboration protocols must specify override authority, audit trails for divergence from AI recommendations, and escalation triggers for when system confidence falls below thresholds.

Strategic talent alignment follows a similar trajectory. At levels 1 and 2, SCP candidate hiring processes screen for ERP proficiency and process discipline, while training focuses on system navigation and transaction accuracy. At levels 3 and 4, hiring processes screen for analytical aptitude and ambiguity tolerance, while training develops programmatic literacy (Python for building analytical processes) and model thinking. This shift occurs because understanding how AI systems generate outputs and where they fail is necessary for SCM professionals (Richey *et al.*, 2023). Finally, at level 5, hiring processes screen for collaborative intelligence, and training develops AI literacy (e.g. model limitations and bias patterns) and sensemaking practices for interpreting outputs. The shift from technical to collaborative emphasis (documented in Section 3.3) reaches its full expression at this level.

In addition to these role-specific changes, advancing beyond level 3 may require structural changes. The role of an SC architect requires decision authority that process-oriented reporting structures cannot accommodate. Organizations must specify where planning authority resides, how AI recommendations flow to decision-makers, and how accountability is assigned when recommendations are accepted or overridden.

4.4 Research agenda

Our findings suggest propositions that respond to calls for socio-technical SCM research on AI interventions (Hendriksen, 2023). Three potential research streams emerge.

Stream 1 examines paradox dynamics and moderators, centered on two points. The first explores the conditions under which the autonomy–ambiguity paradox resolves in favor of augmentation versus overwhelm. Related propositions outline testable relationships among implementation sequencing, decision-rights clarity, and planner outcomes. The second point investigates the conditions under which GenAI supports adaptive responses to genuinely unknowable uncertainty versus predictive responses to quantifiable variation. We connect this to uncertainty-regulation strategies (Sengupta *et al.*, 2025) and propose that augmentation depends on matching the AI system design to the specific type of uncertainty faced by planners.

Stream 2 focuses on intervention effectiveness. Here, the central question is which cognitive load mitigation strategies produce measurable improvement in planner performance and well-being. Propositions in this stream draw on empirical patterns from this study, including the Colgate-Palmolive-related finding that GenAI-powered narrative explanations aid in planner sensemaking. Explainability drives trust formation in human-in-the-loop systems more strongly than predictive performance alone (Sridhar *et al.*, 2026). However, whether this dynamic holds across the longer time horizons and the multi-tier complexity of SCP has not been tested.

Stream 3 extends the analysis to longitudinal outcomes (over 12–36-month horizons). The driving question here concerns how GenAI adoption affects planner decision quality, well-being, and retention over time. Recent evidence linking integrated planning systems to firm performance through disruption (Swink *et al.*, 2025) offers methodological precedent for this stream. To this end, we set forth propositions that distinguish role design effects from workload effects and that specify productivity dip patterns dependent on organizational support.

Table 3 presents illustrative propositions for each stream, with associated methodological requirements. These propositions are not exhaustive; they represent priority directions derived from the transformation framework.

Beyond these propositions, the four transition paradoxes (Figure 3) warrant dedicated investigation. The governance–agility paradox at transition $4 \rightarrow 5$ is particularly underexplored and connects SCP transformation to broader debates on AI ethics and organizational control.

Table 3. Research agenda: streams, propositions, and methodological requirements

Stream	Research question	Propositions	Methodological approach
1. Paradox dynamics and moderators	Under what conditions does the autonomy–ambiguity paradox resolve in favor of augmentation versus overwhelm?	P1a: Organizations automating legacy tasks before expanding AI-generated scenario volume demonstrate lower planner cognitive load than those expanding scenarios first P1b: The productivity J-curve duration correlates inversely with clarity of human–AI decision rights P1c: Organizations implementing GenAI scenario generation without concurrent legacy task automation exhibit higher planner turnover intention than those implementing both simultaneously	Longitudinal designs tracking organizations through the GenAI implementation phases (with cognitive load and role clarity as mediating variables)
	Under what conditions does GenAI support adaptive responses to genuinely unknowable uncertainty versus predictive responses to quantifiable variation?	P1d: GenAI adoption enables simultaneous deployment of predictive, proactive, reactive, and adaptive uncertainty regulation strategies; augmentation outcomes depend on matching the AI system design to the type of uncertainty faced by planners	Comparative case studies across uncertainty contexts Experimental designs manipulating uncertainty type and AI system configuration
2. Intervention effectiveness	Which cognitive load mitigation strategies produce measurable improvements in planner performance and well-being?	P2a: Scenario-filtering heuristics reduce decision time without degrading decision quality when filtering criteria are co-developed with planners P2b: AI-generated narrative explanations reduce cognitive load more effectively than numerical confidence intervals alone P2c: Targeted interventions addressing specific transition paradoxes outperform generic change-management approaches	Experimental and quasi-experimental designs isolating intervention effects and maturity model (Figure 3) as a staging framework for matched comparisons
3. Longitudinal outcomes	How does GenAI adoption affect planner decision quality, well-being, and retention over 12–36-month horizons?	P3a: Planners in augmentation configurations report higher job satisfaction and lower turnover intention than those in overwhelm configurations, controlling for workload P3b: Forecast accuracy and exception resolution time exhibit J-curve patterns with inflection points that are contingent on organizational support mechanisms	Outcome metrics appropriate to SCP contexts (forecast accuracy, exception resolution time, inventory performance), excluding self-reported effectiveness

Source(s): Authors’ own elaboration

5. Conclusion

The SCP role stands at an inflection point. GenAI adoption creates an autonomy–ambiguity paradox that reshapes planner responsibilities without determining outcomes. Whether transformation leads to augmentation or to overwhelm depends on organizational choices beyond the technological capabilities.

This study offers four contributions. First, the empirical analysis identifies two distinct role archetypes (process guardian vs. SC architect), differentiated by mission orientation, tool ecosystem, and competency profile. Second, the transformation framework (Figure 2) models the pathways from GenAI deployment through organizational choice architecture to consequent states. Third, the maturity model (Figure 3) operationalizes this transformation framework into diagnostic stages with transition-specific paradoxes and resolution requirements. Fourth, the methodological firewall separating classification data from analysis data addresses validity concerns in qualitative technology-adoption research and, as an additional finding, highlights a lag in job specification documentation.

The study's empirical base remains narrow. Four adopter cases and eight non-adopter cases across four industry sectors constrain its generalizability, and the overwhelm pathway is theoretically derived from role-conflict literature and extrapolated from the transformation framework, as no observed cases occupy the pathway directly. Longitudinal tracking of organizations through their GenAI implementation would test the propositions in Table 3 and reveal whether the claimed productivity dip patterns and path dependencies operate as the framework predicts.

Notes

1. This paper defines generative AI (GenAI) as computational systems that produce novel outputs (e.g. forecasts, recommendations, narratives, structured decision options) from learned data distributions, distinct from discriminative or optimization models that classify, rank, or select among existing patterns. Agentic AI systems that autonomously plan and execute multi-step workflows qualify as GenAI deployments when their reasoning and instruction-generation engine uses a generative foundation model; the coordination layer alone does not alter that classification. Both agentic and generative AI build on generative foundations, but the two modes produce qualitatively distinct organizational consequences (Chalmers *et al.*, 2026). Section 2.3 specifies how this definition was operationalized for case classification.
2. This does not imply that these organizations utilize no advanced analytics or AI in other parts of their business (e.g. marketing, finance), only that our classification was strictly limited to documented applications of GenAI within the scope of SCP, according to the specified sources.

Supplementary material

The supplementary material for this article can be found online.

References

Acemoglu, D. and Restrepo, P. (2019), "Automation and new tasks: how technology displaces and reinstates labor", *Journal of Economic Perspectives*, Vol. 33 No. 2, pp. 3-30, doi: 10.1257/jep.33.2.3.

Aghion, P. and Tirole, J. (1997), "Formal and real authority in organizations", *Journal of Political Economy*, Vol. 105 No. 1, pp. 1-29, doi: 10.1086/262063.

Ansari, A.F., Stella, L., Turkmen, C., Zhang, X., Mercado, P., Shen, H., Shchur, O., Rangapuram, S.S., Arango, S.P., Kapoor, S. and Zschiegner, J. (2024), "Chronos: learning the language of time series", *Transactions of Machine Learning Research*, pp. 1-46.

Atalay, E., Phongthiengtham, P., Sotelo, S. and Tannenbaum, D. (2020), "The evolution of work in the United States", *American Economic Journal: Applied Economics*, Vol. 12 No. 2, pp. 1-34, doi: 10.1257/app.20190070.

Bettis, R.A. (2017), “Organizationally intractable decision problems and the intellectual virtues of heuristics”, *Journal of Management*, Vol. 43 No. 8, pp. 2620-2637, doi: 10.1177/0149206316679253.

Boyer, K.K. and Swink, M.L. (2008), “Empirical elephants—why multiple methods are essential to quality research in operations and supply chain management”, *Journal of Operations Management*, Vol. 26 No. 3, pp. 338-344, doi: 10.1016/j.jom.2008.03.002.

Brynjolfsson, E., Rock, D. and Syverson, C. (2021), “The productivity J-curve: how intangibles complement general purpose technologies”, *American Economic Journal: Macroeconomics*, Vol. 13 No. 1, pp. 333-372, doi: 10.1257/mac.20180386.

Brynjolfsson, E., Li, D. and Raymond, L. (2025), “Generative AI at work”, *The Quarterly Journal of Economics*.

Canciglieri, A.B., Seelent, J.F.C., Kai, D.A., Franco, C.W. and Benitez, G.B. (2026), “Adaptive leadership in artificial intelligence solutions in operations and supply chain management projects: the Robert Bosch GmbH case”, *International Journal of Physical Distribution and Logistics Management*, Vol. ahead-of-print No. ahead-of-print, pp. 1-22, doi: 10.1108/ijpdlm-12-2024-0497.

Chalmers, D., Hunt, R.R., Pachidi, S., Potočnik, K. and Townsend, D. (2026), “The acceleration of artificial intelligence: rethinking organization and work in an era of rapid technological change”, *Journal of Management Studies*, Vol. 63 No. 2, pp. 285-314, doi: 10.1111/joms.70063.

Christopher, M. (2016), *Logistics and Supply Chain Management*, Pearson, UK.

Ciceri, C., Caniato, F. and Moretto, A. (2026), “Integrating artificial intelligence capabilities in supply chain cyber risk management”, *International Journal of Physical Distribution and Logistics Management*, pp. 1-38, doi: 10.1108/ijpdlm-01-2025-0028.

Dubey, R., Gunasekaran, A. and Papadopoulos, T. (2024), “Benchmarking operations and supply chain management practices using Generative AI: towards a theoretical framework”, *Transportation Research Part E: Logistics and Transportation Review*, Vol. 189, 103689, doi: 10.1016/j.tre.2024.103689.

Eisenhardt, K.M. (1989), “Building theories from case study research”, *Academy of Management Review*, Vol. 14 No. 4, pp. 532-550, doi: 10.2307/258557.

Flöthmann, C., Hoberg, K. and Wieland, A. (2018), “Competency requirements of supply chain planners and analysts and personal preferences of hiring managers”, *Supply Chain Management: An International Journal*, Vol. 23 No. 6, pp. 480-499.

Gioia, D.A., Corley, K.G. and Hamilton, A.L. (2013), “Seeking qualitative rigor in inductive research: notes on the Gioia methodology”, *Organizational Research Methods*, Vol. 16 No. 1, pp. 15-31, doi: 10.1177/1094428112452151.

Gligor, D.M., Holcomb, M.C. and Stank, T.P. (2013), “A multidisciplinary approach to supply chain agility: conceptualization and scale development”, *Journal of Business Logistics*, Vol. 34 No. 2, pp. 94-108, doi: 10.1111/jbl.12012.

Hendriksen, C. (2023), “Artificial intelligence for supply chain management: disruptive innovation or innovative disruption?”, *Journal of Supply Chain Management*, Vol. 59 No. 3, pp. 65-76, doi: 10.1111/jscm.12304.

Huynh, M.-T. (2024), “Using generative AI as decision-support tools: unraveling users’ trust and AI appreciation”, *Journal of Decision Systems*, pp. 1-32, doi: 10.1080/12460125.2024.2428166.

Ivanov, D., Dolgui, A. and Sokolov, B. (2019), “The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics”, *International Journal of Production Research*, Vol. 57 No. 3, pp. 829-846, doi: 10.1080/00207543.2018.1488086.

Jackson, I., Ivanov, D., Dolgui, A. and Namdar, J. (2024), “Generative artificial intelligence in supply chain and operations management: a capability-based framework for analysis and implementation”, *International Journal of Production Research*, Vol. 62 No. 17, pp. 6120-6145, doi: 10.1080/00207543.2024.2309309.

Jonsson, P. and Holmström, J. (2016), “Future of supply chain planning: closing the gaps between practice and promise”, *International Journal of Physical Distribution and Logistics Management*, Vol. 46 No. 1, pp. 62-81, doi: 10.1108/ijpdlm-05-2015-0137.

Kulik, C.T. and Perry, E.L. (2023), *Human Resources for the Non-HR Manager*, Taylor & Francis, New York.

Lauchlan, S. (2025), "How agentic AI will help Colgate-Palmolive in its pursuit of 'perceivable superiority'...with human intelligence factored in", *Diginomica*, available at: https://diginomica.com/how-agentic-ai-will-help-colgate-palmolive-its-pursuit-perceivable-superioritywith-human (accessed 4 December 2025).

Levinthal, D.A. and Wu, B. (2025), "Resource redeployment and the pursuit of the new best use: economic logic and organizational challenges", *Strategy Science*, Vol. 10 No. 1, pp. 32-47, doi: 10.1287/stsc.2022.0105.

Li, L., Liu, Y., Jin, Y., Cheng, T.C.E. and Zhang, Q. (2024), "Generative AI-enabled supply chain management: the critical role of coordination and dynamism", *International Journal of Production Economics*, Vol. 277, 109388, doi: 10.1016/j.ijpe.2024.109388.

Mandal, S., Babu, A.S., Dubey, R.K., M, S. and M, S. (2026), "From infrastructure to insight: a systemic analysis of AI adoption barriers in supply chain forecasting", *International Journal of Productivity and Performance Management*, pp. 1-31, doi: 10.1108/ijppm-05-2025-0345.

Parker, S.K. and Grote, G. (2022), "Automation, algorithms, and beyond: why work design matters more than ever in a digital world", *Applied Psychology*, Vol. 71 No. 4, pp. 1171-1204, doi: 10.1111/apps.12241.

Pentland, B.T. and Feldman, M.S. (2005), "Organizational routines as a unit of analysis", *Industrial and Corporate Change*, Vol. 14 No. 5, pp. 793-815, doi: 10.1093/icc/dth070.

Puranam, P. (2021), "Human–AI collaborative decision-making as an organization design problem", *Journal of Organization Design*, Vol. 10 No. 2, pp. 75-80, doi: 10.1007/s41469-021-00095-2.

Raisch, S. and Krakowski, S. (2021), "Artificial intelligence and management: the automation–augmentation paradox", *Academy of Management Review*, Vol. 46 No. 1, pp. 192-210, doi: 10.5465/amr.2018.0072.

Richey, R.G., Chowdhury, S., Davis-Sramek, B., Giannakis, M. and Dwivedi, Y.K. (2023), "Artificial intelligence in logistics and supply chain management: a primer and roadmap for research", *Journal of Business Logistics*, Vol. 44 No. 4, pp. 532-549, doi: 10.1111/jbl.12364.

Seeber, I., Bittner, E., Briggs, R.O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A.B., Oeste-Reiß, S., Randrup, N., Schwabe, G. and Söllner, M. (2020), "Machines as teammates: a research agenda on AI in team collaboration", *Information and Management*, Vol. 57 No. 2, 103174, doi: 10.1016/j.im.2019.103174.

Sengupta, S., Jonsson, P., Dreyer, H.C., Kaipia, R. and Choi, T.Y. (2025), "Uncertainty regulation and adaptable supply chain planning", *International Journal of Operations and Production Management*, Vol. 45 No. 11, pp. 1861-1883, doi: 10.1108/ijopm-07-2025-0645.

Shrestha, Y.R., Ben-Menahem, S.M. and von Krogh, G. (2019), "Organizational decision-making structures in the age of artificial intelligence", *California Management Review*, Vol. 61 No. 4, pp. 66-83, doi: 10.1177/0008125619862257.

Smith, W.K. and Lewis, M.W. (2011), "Toward a theory of paradox: a dynamic equilibrium model of organizing", *Academy of management Review*, Vol. 36 No. 2, pp. 381-403, doi: 10.5465/amr.2009.0223.

Sridhar, S., Baskar, P., Grimes, J. and Sampathkumar, A. (2026), "A comprehensive framework for human - AI collaborative decision making in intelligent retail environments", *Expert Systems with Applications*, Vol. 299, 130013, doi: 10.1016/j.eswa.2025.130013.

Strauss, A. and Corbin, J. (1998), *Basics of Qualitative Research Techniques*, Sage Publications, Thousand Oaks, CA.

Swink, M., Kovach, J.J. and Roh, J. (2025), "Inventory and supply chain planning systems as drivers of supply chain resilience: analyses of firm performance through the COVID-19 pandemic", *Production and Operations Management*, Vol. 34 No. 8, pp. 2486-2505, doi: 10.1177/10591478251320528.

Sydow, J., Schreyögg, G. and Koch, J. (2009), "Organizational path dependence: opening the black box", *Academy of Management Review*, Vol. 34 No. 4, pp. 689-709, doi: 10.5465/amr.2009.44885978.

Tan, K.H., Zhan, Y., Ji, G., Ye, F. and Chang, C. (2015), "Harvesting big data to enhance supply chain innovation capabilities: an analytic infrastructure based on deduction graph", *International Journal of Production Economics*, Vol. 165, pp. 223-233, doi: 10.1016/j.ijpe.2014.12.034.

Thaler, R.H., Sunstein, C.R. and Balz, J.P. (2013), "Choice architecture", in Shafir, E. (Ed.), *The Behavioral Foundations of Public Policy*, Princeton University Press, Princeton, NJ, pp. 428-439.

Trkman, P., McCormack, K., de Oliveira, M.P.V. and Ladeira, M.B. (2010), "The impact of business analytics on supply chain performance", *Decision Support Systems*, Vol. 49 No. 3, pp. 318-327, doi: 10.1016/j.dss.2010.03.007.

Tubre, T.C. and Collins, J.M. (2000), "Jackson and schuler (1985) revisited: a meta-analysis of the relationships between role ambiguity, role conflict, and job performance", *Journal of Management*, Vol. 26 No. 1, pp. 155-169, doi: 10.1177/014920630002600104.

Wamba, S., Queiroz, M.M., Chiappetta Jabbour, C.J. and Shi, C. (2023), "Are both generative AI and ChatGPT game changers for 21st-Century operations and supply chain excellence?", *International Journal of Production Economics*, Vol. 265, 109015, doi: 10.1016/j.ijpe.2023.109015.

Wamba, S.F., Guthrie, C., Queiroz, M.M. and Minner, S. (2024), "ChatGPT and generative artificial intelligence: an exploratory study of key benefits and challenges in operations and supply chain management", *International Journal of Production Research*, Vol. 62 No. 16, pp. 5676-5696, doi: 10.1080/00207543.2023.2294116.

Wieland, A. (2021), "Dancing the supply chain: toward transformative supply chain management", *Journal of Supply Chain Management*, Vol. 57 No. 1, pp. 58-73, doi: 10.1111/jscm.12248.

Wilson, H.J. and Daugherty, P.R. (2018), "Collaborative intelligence: humans and AI are joining forces", *Harvard Business Review*, Vol. 96 No. 4, pp. 114-123.

Xiang, Y., Vélez, N. and Gershman, S.J. (2024), "Optimizing competence in the service of collaboration", *Cognitive Psychology*, Vol. 150, 101653, doi: 10.1016/j.cogpsych.2024.101653.

Corresponding author
Patrik Jonsson can be contacted at: patrik.jonsson@chalmers.se