



Exploring the Impact of Agentic AI in the CPG Supply Chain

In the space of a year, AI agents have moved from pitch deck to product across the major supply chain planning platforms. Gartner forecasts that supply chain management software with agentic AI capabilities will grow from less than \$2 billion in 2025 to \$53 billion by 2030. NVIDIA's 2026 State of AI in Retail and CPG survey found that 47% of respondents are using or assessing agentic AI, with 20% saying agents are already active in their organizations.

Agentic AI is already part of supply chain planning. The more important question is where CPG companies should trust agents to make decisions. For most organizations, the near-term opportunity for agentic AI in the CPG supply chain isn't fully autonomous planning. It's applying governed autonomy to specific workflows where agents can improve speed and decision quality while keeping people accountable for higher-risk decisions.

What "Agentic" Actually Changes

The difference that matters is autonomy.

Traditional [supply chain AI](#) sensed demand, optimized inventory, and recommended actions, then waited for a planner to act. Agentic AI goes a step further. It pursues a goal, reasons across systems, recommends or takes action through approved workflows, and captures outcomes so the process can improve over time.

That doesn't replace forecasting or optimization. Predictive models remain the foundation of supply chain planning. Agentic AI changes what happens after the forecast is produced. Rather than asking planners to review every exception, agents determine which issues deserve attention, evaluate response options, update systems where appropriate, and identify [when a human needs to step in](#).

Consider an inventory agent with a standing objective of maintaining service levels above 98% while minimizing excess inventory. An inventory agent can monitor demand, supplier performance, inventory positions, and transportation constraints. When disruption occurs, it identifies likely stockouts, evaluates response options, and either recommends or executes approved actions within established decision rights.

What happens next depends on the organization's governance model. The agent may evaluate options such as DC rebalancing, expedited replenishment, alternate sourcing, order allocation, or safety-stock adjustment. Depending on its assigned decision rights, it can recommend an action, execute pre-approved steps within thresholds, or escalate higher-impact decisions to a planner.

From Prediction to Action: How the Two Build on Each Other

Predictive AI

Senses and recommends, then stops

Senses demand and signals



Produces a forecast



Recommends an action



Waits for a person to act

Agentic AI

Pursues a goal and acts on it

Pursues a standing goal



Reasons across systems



Take actions



Learns and adjusts

[Whether an organization is ready](#) for that level of autonomy depends less on the technology use than on the fundamentals: data quality, system integration, controls, auditability, and risk appetite. Understanding what agents can do is only the starting point. The next challenge is deciding where organizations should trust them to act and where human oversight still delivers better outcomes.

Where Agents Can Create Value First

The pull toward agents is strongest where the payoff is easy to prove. In CPG, the clearest early returns are in supply chain operations - the places where faster exception handling, better availability, and lower inventory show up directly in the numbers. The best approach is straightforward: start with a problem tied to a specific metric, demonstrate the value, then gradually expand the agent's decision authority.

A practical CPG roadmap can be organized by three levels of autonomy:

Advisor: recommends, human decides

At this stage, agents improve decision quality while planners retain authority.

- **Demand sensing:** An agent reads point-of-sale, syndicated, order, inventory, and external signals, [such as weather](#), to identify forecasts that require planner attention.
- **Availability gap and root-cause detection:** An agent scans shelf, inventory, and DC data each day, surfaces the most urgent out-of-stock risks, and traces what caused them. Kraft Heinz Away From Home reported that [an AI agent identified a risk](#) of 6,200 cases (*approximately \$300,000 monthly*) across five SKUs, allowing the team to work with a partner before the issue grew.
- **Stockout resolution:** An agentic workflow can identify shortage or unfulfilled demand risks, explain the constraints driving them, and guide planners toward mitigation actions such as scenario planning, stock transfer requisitions, purchase requisitions, deployment changes, or prioritized order confirmation. SAP's Planning Assistant, supported by [Joule agents in SAP IBP and SAP S/4HANA Backorder Processing](#), is one example of this type of guided workflow.
- **Constraint-aware planning:** An agent evaluates supply constraints and recommends the next-best action using current planning context. [Kinaxis positions its Maestro Agents](#) as context-aware digital coworkers embedded in live planning environments to "plan with you, not for you."

Bounded actor: acts within limits

As confidence grows, organizations can allow agents to execute low-risk decisions within predefined guardrails.

- **Replenishment and network rebalancing:** An agent reviews on-hand inventory, on-order supply, demand risk, DC inventory, and service-level targets, then recommends or executes a transfer, order change, or expedite inside approved thresholds.

- **Supplier and transportation exception triage:** An agent gathers the latest purchase order, ASN, carrier, lead-time, and risk signals, prepares the best mitigation path, and escalates only when cost, service, or policy thresholds are exceeded.
- **Guided execution:** [SAP says](#) it's delivering more than 60 purpose-built supply chain agents designed to sense events, analyze impact, and take guided action within defined business guardrails while keeping people in control.

Autonomous workflow: future state, not starting point

Fully autonomous supply chains remain more of a destination than today's reality.

Vendors increasingly describe autonomous planning as the long-term vision. [SAP](#), for example, describes its roadmap as the "autonomous supply chain." Today, however, most successful implementations focus on narrow, repetitive decisions where the operational risk is low and incorrect actions are easily reversed. Examples include automatically resolving routine planning exceptions or updating planning parameters within approved limits.

These examples show the current market pattern. The strongest near-term agents aren't generic chatbots. They're embedded in planning, replenishment, retail data, and execution workflows, connected to the data the decision depends on, and governed by clear guardrails. Their value comes from shifting routine analysis from planners to the system while keeping accountability visible.

Three Levels of Agent Autonomy

The same agent, the same learning mindset. What changes is how much it is allowed to do.



Advisor

Recommends,
human decides

AGENT: Highlights risk, explains why, and suggests a fix.

HUMAN: Makes the call and executes.



Bounded

Acts within limits

AGENT: Handles low-risk scenarios in a row, escalates the rest.

HUMAN: Sets thresholds and reviews escalations.



Autonomous

Acts end to end

AGENT: Handles end-to-end, plans the order, then adjusts.

HUMAN: Sets the goal and the guardrails.

Common Starting Point

Autonomy expands as data, trust, and governance mature.

Capability Gets You an Agent. Trust Gets You Adoption.

Here is the part that should shape how CPG companies invest: the main reason AI projects fail is less often model capability and more often operational fit. *Do planners trust those recommendations enough to act on them?*

MIT's 2025 GenAI Divide report found that [95% of enterprise generative AI pilots](#) produced no measurable financial impact, despite significant investment. Gartner makes a similar point for agentic AI, predicting that more than [40% of agentic AI projects](#) will be canceled by the end of 2027 because of escalating costs, unclear value, or inadequate risk controls.

Those statistics reflect a simple problem. If planners can't understand or trust an agent's recommendations, they'll override them. Adoption depends as much on changing decision-making as improving the technology.

That shift also changes the planner's role. Rather than spending time gathering information, planners increasingly evaluate an agent's recommendations, recognize when supporting data is incomplete, determine when recommendations conflict with business priorities, and know when to escalate. The [World Economic Forum](#) expects 39% of workers' core skills to change by 2030. In supply chain, that change will show up as AI literacy, exception judgment, cross-functional decision-making, and comfort working through human-machine workflows.



Designing Agents People Will Actually Use

In practice, governance means defining decision rights, approval thresholds, escalation paths, audit trails, and outcome measurement before expanding an agent's authority. The companies that pull ahead will treat adoption as the design problem, not something to fix after launch.

A few questions separate an agent that earns its keep from one that gets ignored:

- What business metric is the agent expected to move? Service level, forecast-value-add, inventory turns, expedite cost, case fill, planner productivity, or lost sales?
- Can a planner see and challenge the agent's reasoning, assumptions, and source data?
- Which decisions can the agent make on its own? Which require approval, and which must always escalate?
- Are the actions reversible, auditable, and tied to clear accountability?
- When the agent and planner disagree, who makes the final decision? How is that outcome captured for improvement?

A practical adoption path expands agent autonomy gradually. Start with advisor use cases where the agent explains risk and suggests options. Move to bounded action once data quality and controls are reliable, with clear escalation rules in place. Reserve autonomous workflows for high-volume decisions with limited risk, where the available actions are well defined and the impact of an incorrect decision is manageable.

Moving Forward

Agentic AI will change supply chain planning, but adoption will happen one workflow at a time. Organizations that succeed will introduce autonomy where it improves business outcomes, build trust through governance, and expand decision authority as confidence grows.

Clarkston helps CPG organizations turn agentic AI from a promising pilot into something teams use every day, aligning the technology with the processes, ownership, controls, and trust required to make it pay off. If you're evaluating where agentic AI fits within your supply chain strategy, let's talk.

ABOUT CLARKSTON CONSULTING

Businesses across the consumer products industry partner with Clarkston Consulting to enhance strategic decision-making, improve operational efficiencies, implement new technologies, and promote business growth and market diversification. Leveraging deep functional and industry expertise, our people discover, design, and deliver solutions that fit your business, your goals, and your future. At Clarkston, your purpose is our purpose.

For more information about how we can help your company, [please contact us](#).

