AI & PLANNING SYSTEMS

Inside the AI Circus: When Consulting Hype Replaces Thinking

Christian Kroschl · 24 November 2025

Article on kroschl.com · Read on LinkedIn

Every week brings another "AI-enabled planning" workshop. Same setup. Same buzzwords. Same noise. AI Agents. Self-Learning Ecosystems. Autonomous Supply Chains.

And every time, I leave with the same question:

> *"Does anyone in the room actually understand how planning works?"*

The New Consulting Religion

AI has become a belief system. Slide after slide promises "self-driving supply chains" — as if data, processes, and systems would somehow start running themselves. Yet when you ask the simplest question — *what exactly is autonomous here?* — the conversation collapses.

I've been through a two-hour session with a large life sciences client last week. Dozens of people. Dozens of slides. Not a single definition of what "AI in planning" actually meant. We spent more time talking about *ecosystems* and *customer-centric visions* than about lead times, capacities, or buffers — the stuff that actually makes a supply chain run.

The Echo Chamber

Too many strategy calls have turned into echo chambers. Everyone repeats the same words: *alignment, orchestration, end-to-end visibility.* No one asks the only question that matters: *what are we solving?*

One senior partner closed the session by saying, *"Let's continue this journey towards something tangible."* After nearly two hours of talking — still no tangible definition, no ownership, no outcome. That's the consulting version of white noise.

The Only Moment of Clarity

When I explained that AI in planning is a **stochastic problem** — driven by uncertainty, probability, and variance — the room went quiet. Because that's where the hype meets reality. Planning isn't an app you download. Neural networks don't "self-learn" without structured, contextualized data. And large language models don't plan — they predict words.

A life sciences supply chain isn't a chatbot. It's a regulated, validated system that requires traceability, accuracy, and control. Planners won't be replaced by LLMs; they'll be **augmented by better stochastic models** — if we get the basics right.

The Blind Teaching the Blind

Too many AI conversations suffer from a simple problem: **the blind are teaching the blind.**

People without statistical grounding explain "AI strategy" to leaders who have never built or operated a real planning model. It's not a talent issue. It's not a motivation issue. It's a **competence gap wrapped in confident language.**

If we cannot distinguish machine learning from LLMs — or variance from noise — then our AI vision is built on metaphors, not models.

The Data Mirage

I spent 2.5 hours two weeks ago with a senior data lead from a major life sciences company. His story is the same one I hear everywhere:

No single source of truth. No harmonized master data. No end-to-end data model. Factories planning in Excel because they don't even run a standardized ERP backbone.

Everyone talks about *Snowflake, Dataland, Data Lakes* — but what they have are isolated, use-case-specific data pots without a shared semantic layer.

From a data-architecture perspective, the problem is simple:

> ***You cannot build a data lake when you don't have a river.***

SAP HANA is widely used and could serve as a strong transactional foundation — complete with document flow, referential integrity and process lineage — but often not used as such. Instead, companies replicate fragmented extracts into yet another data platform, hoping AI will somehow fix structural inconsistency.

But AI cannot learn from noise.

In statistics, it's called a *broken data-generating process* — non-stationary, inconsistent, impossible to model, impossible to trust. In enterprise architecture, it is called what it is: **an absence of end-to-end process logic.**

Processes sound boring. But processes *generate* data. Input → Process → Output. Output becomes the input of the next step. ERP systems were designed around this simple principle.

If there is no harmonized process, there is no harmonized data. If there is no harmonized data, there is no AI — only the illusion of it.

But let's keep inflating the AI bubble — until the lack of foundations finally makes it burst.

The Ownership Void

Data issues are not only technical — they are structural. Ownership is highly fragmented. Data teams may manage material, vendor and customer data, but **manufacturing data falls through the cracks**.

Who owns the accuracy of BOMs? Who validates recipes and routings? Who ensures purchasing info records reflect the right transport lanes or sourcing priorities?

In most companies, the truthful answer is: *no one end-to-end.*

Why? Because the organization still doesn't think in value chain terms. Supply chain may have moved out of the back office, but **end-to-end data ownership, process accountability and planning logic** are still treated as optional.

The principles are simple:

Data follows process. Process precedes data modelling. Planning logic defines the data concept. Organization follows process — not the other way around.

Yet these fundamentals remain widely misunderstood and rarely followed.

Why It Keeps Failing

AI in supply chain doesn't fail because of technology. It fails because:

Leaders chase headlines, not foundations. Consultants sell vision, not design. Teams talk in circles instead of defining ownership.

It's not an *AI problem* — it's a *thinking problem.*

What Real AI Looks Like

Real AI in planning is boring. It's about machine learning for forecast patterns. It's about optimization models that balance constraints. It's about prescriptive logic that supports decisions in real time. And all of that depends on clean data, stable parameters, and closed feedback loops — not sticky notes and slogans.

What Leaders Must Fix First

To avoid inflating the AI bubble any further, leaders need to anchor execution in four fundamentals:

Define end-to-end process ownership – without clear accountability, data and planning will always

drift. **Establish a value-chain data model** – including shared dimensions, golden records, and explicit data contracts. **Stabilize parameters and feedback loops** – AI needs consistent signals, not random noise. **Sequence transformation correctly** – process → planning logic → data → AI. Never the other way around.

These are not optional. They are prerequisites.

Strengthening the Technical Foundation

Modern ML systems depend on **feature stability**, **consistent semantics**, and **reliable feedback loops**. Without them, every model will suffer from concept drift, misclassification, or brittle predictions.

Data contracts — the explicit agreements defining structure, semantics and lineage between producing and consuming systems — are essential to prevent these failures. Most companies still lack them entirely.

Final Thought

> *The more we talk about "autonomous supply chains," the less autonomous our thinking becomes. Supply chains don't need another vision deck — they need people who can still tell signal from noise.*

#SupplyChain #AI #MachineLearning #LLM #APS #ExecutionMatters #Resilience #PlanningReality #LifeSciences #EY

What "agentic AI" really is – beyond the buzzword

In research, an *agent* is a system that perceives its environment, pursues a goal and takes actions to move towards that goal. Agentic AI is simply this idea applied to today's models: a model (often an LLM) that can plan over multiple steps, call tools and APIs, observe the results, update its context, and then decide what to do next.

In practice, most so-called "AI agents" today are nothing more than a loop: understand the goal, plan the next step, call a tool, read the output, repeat. It's powerful – but only if the tools, data and processes behind it are solid.

In supply chain terms, a realistic example is a planning copilot that monitors forecast errors, OTIF and inventory, detects an exception, pulls the right data from ERP and APS, runs a few scenarios and then prepares a recommendation for the planner. The system is *agentic* – but it is not a magic autonomous supply chain. It still depends on clear processes, a robust data model and explicit decision policies.

As of 2025, this is where we really are: AI agents work well for narrow, well-bounded tasks with good data. They struggle with long-horizon autonomy across fragmented systems, weak master data and unclear ownership. We should absolutely experiment with agents – but we should be honest that they are orchestrators on top of our architecture, not a shortcut around missing end-to-end thinking.