


BEYOND THE CO-PILOT

From Co-Pilot to Colleague — The Operating Model Shift Nobody Is Planning For

A Practitioner's Perspective on Agentic AI and Organisational Readiness

Giovanni Leonardi

October 2024



From Co-Pilot to Colleague — The Operating Model Shift Nobody Is Planning For

Leadership · No. 93

Originally written October 2024 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"Most organisations have no answer to that question — not because they have considered it and found it difficult, but because they have not yet recognised it as a question that needs asking."

The Comfortable Illusion of the Co-Pilot

For the past eighteen months, organisations have settled into a comforting narrative about artificial intelligence in the workplace. AI is a co-pilot — a helpful assistant that augments human capability, drafts the email, summarises the meeting, suggests the code. The human remains in command. The tool does what it is told.

This narrative is already outdated.

The shift from co-pilot to colleague — from AI that assists to AI that acts — is happening faster than most operating models can absorb. Agentic AI systems that set their own sub-goals, execute multi-step tasks, and make consequential decisions with minimal human oversight are no longer research prototypes. They are entering production environments. And the governance, leadership, and organisational design frameworks that were built for a world of human-directed tools are simply not equipped for a world of machine-initiated action.

Guardrails Built for the Wrong Problem

The pattern I have observed across sectors is remarkably consistent. Organisations that adopted generative AI early moved quickly to establish acceptable-use policies, data governance frameworks, and responsible-AI principles. These were sensible and necessary steps — but they were designed for a specific model of interaction: a human asks, a machine responds, the human decides.

Agentic AI breaks this model at its foundation. When an AI system autonomously triages customer complaints, re-prioritises a work queue, or triggers a procurement workflow based on its own assessment of conditions, the locus of decision-making has shifted. The question is no longer "Is this tool being used responsibly?" but "Who is accountable when the agent acts on its own judgement?"

Most organisations have no answer to that question — not because they have considered it and found it difficult, but because they have not yet recognised it as a question that needs asking.

The Operating Model Gap

The real challenge is not technological. The agentic frameworks are maturing rapidly, and the engineering community is solving the technical problems of tool use, memory, and planning with characteristic speed. The challenge is organisational.

Operating models define who decides, who is accountable, and how work flows through an organisation. When a new class of actor enters that system — one that is neither employee nor tool — the operating model must evolve to accommodate it. Almost none have.

Consider the implications across three dimensions:

- **Accountability structures.** Traditional operating models assign accountability to named individuals within defined roles. An AI agent occupies no role, holds no accountability, and cannot be disciplined or redirected through normal management channels. The accountability gap is not a technicality — it is a governance vacuum.
- **Decision rights.** Organisations invest considerable effort in defining decision rights — who can approve what, at what threshold, with what oversight. Agentic systems that make or materially influence decisions sit outside these frameworks entirely. They operate in the spaces between defined authorities, often at a speed that precludes the human review the framework assumed.
- **Workflow design.** Most workflows are designed around human handoffs, human judgement points, and human escalation paths. Inserting an autonomous agent into these workflows does not simply accelerate them — it fundamentally changes their character. The workflow becomes a hybrid system, and hybrid systems have failure modes that neither purely human nor purely automated processes exhibit.

Leadership Models That Do Not Yet Exist

In my experience, the leadership challenge is the most underestimated dimension of this shift. Managing a team of humans who use AI tools is a familiar extension of existing management practice. Managing a portfolio of work where some tasks are performed by humans, some by agents, and some by human-agent partnerships is something qualitatively different.

The leader in this environment needs to understand not just what the agent can do, but how it reasons, where its confidence is warranted, and where its judgement should be overridden. This is not a technical skill — it is a new form of operational literacy that sits somewhere between management and engineering, and for which no development

pathway currently exists.

What This Demands

The temptation — and I see it repeatedly — is to treat agentic AI as the next iteration of automation and to govern it with the same frameworks. This is a category error. Automation executes predefined logic. Agentic AI exercises judgement. The governance requirements are fundamentally different.

What is needed, and what almost no organisation has begun to build, is a governance model designed for supervised autonomy — one that defines the boundaries within which an agent may act independently, the triggers for human escalation, the audit trail for agent-initiated decisions, and the accountability framework that connects agent actions to human responsibility.

The organisations that will navigate this transition successfully are not those with the best AI technology, but those that recognise the operating model must change before the technology forces it to. This is not a technology programme. It is an organisational design challenge, and it requires leadership attention at the most senior level. The organisations that begin this work now — before the technology forces the conversation — will have a significant advantage over those that wait for the first serious failure to make the case for them.