pwc | Microsoft

Unlocking Tomorrow:

A playbook that explores the power of agentic AI for consumer markets

Table of contents

Executive Summary

Agentic AI — autonomous agent systems that can understand intent and plan, execute, orchestrate and adapt activities to deliver complex outcomes, involving humans where judgements matter — is reshaping how consumers discover products, how brands reach them, and how enterprises operate. AI agents that shop, compare, and transact on behalf of consumers are forecast to **drive up to 15%** of European e-commerce spending by 2030, with adoption spreading at **four times the pace** of traditional e-commerce. Across the enterprise, agents are beginning to reshape operational processes from product development to supply chain orchestration and customer service. Retail and consumer goods companies sit at the centre of both shifts: the commercial model, the value chain, and the workforce are all being reshaped at once.

Most companies recognise the imperative, but few are capturing the value. **PwC's 2026 AI Performance Study** shows that the top 20% of firms now generate 74% of all AI-driven returns, outperforming the rest by a factor of 7.2. The gap isn't about who started first. It's about who scales — and scaling requires more than deploying agents. It requires redesigning how the business operates: end-to-end processes, data foundations, governance, and the way people work. Along with robust technology and solid organisational structures, these are the bedrocks that make an enterprise agent-ready.

the top **20%** of firms now generate

74%

of all AI-driven returns

Source: PwC's 2026 AI Performance Study

This article maps the path from experimentation to enterprise impact. What's included:

- Eight value loops where agentic AI creates the greatest returns in retail and consumer goods
- Why most firms stall at the pilot stage — and why workforce and operating-model redesign are the binding constraints to scale
- How trust and responsible AI become the license to scale, not a compliance overhead
- The three-layer architecture that agent-ready enterprises are building
- Five priorities for leadership teams
- How PwC and Microsoft can support the journey and help you scale agentic AI

01

Agentic AI is reshaping retail and consumer goods

The retail and consumer goods landscape is changing radically. Supply chain shifts, tariff dynamics, and evolving consumer confidence are testing traditional operating models — and creating a defining opportunity for those ready to reinvent. Agentic AI is emerging as a catalyst for that reinvention: not another technology layer, but a new way to embed intelligence, speed, and adaptability into the core of the business.

Most organisations have built their digital capabilities over many years, but often they aren't yet fully connected. Agentic AI is beginning to change the equation. Agents can observe, decide, and act across workflows, embedding intelligence directly into operations rather than layering it on top. AI agents can connect what was disconnected and and turn data into outcomes in real time. These impacts are beginning to transform internal company operations across many industries.

In the retail and consumer goods industry, external changes in customer behaviour are also already starting to have significant impacts. The rise of what's known as "agentic commerce" describes a world where AI agents are used to discover and buy goods. By 2030, **33% of consumers** expect to make purchases through AI agents. Agent-driven transactions could route **more than €100 billion** through non-human decision-makers across Europe, with adoption spreading at up to four times the pace of traditional e-commerce in its early years. And **20% of retailers** have already deployed agents along the value chain.

€100 bn

agent-driven transactions could route more than €100 billion through non-human decision-makers across Europe

33%

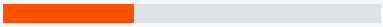

of consumers expect to make purchases through AI agents **by 2030**

20%

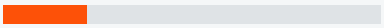

of retailers have already deployed agents along the value chain

Source: The agentic AI revolution in retail

This isn't a front-end story. It is a redesign of the commercial model.

For decades digital commerce followed a predictable path: the customer searched, compared, and checked out. Brands competed for visibility at each step. That sequence is being rewritten. Discovery is now beginning to happen upstream, taking place before customers even reach a traditional digital shelf. In the future, the new flow may happen smoothly behind the scenes: a customer expresses intent, an agent synthesises options, routes demand, and the transaction executes, often without a single page view.

What this means for retail

The implications for retail are fundamental. Discoverability is no longer a merchandising problem. It is a data, trust, and architecture challenge. Agents evaluate structured data, fulfilment reliability, pricing logic, and trust signals. Being machine-readable matters as much as being visually compelling. Most e-commerce systems weren't built for queries like "Is this the healthier option?" or "Am I getting the most value for my money?". But that's where some buying decisions now start.

The impact varies by category. For commodities (screws, batteries, basic household supplies), agents accelerate commoditisation: brand matters less when the agent simply fetches whatever meets the specification at the best price. For differentiated categories (luxury, beauty, lifestyle), the calculus is different but the dual-track challenge is real. Brands need rich visual experiences for human shoppers and machine-readable structured data for their agents, both simultaneously. The companies that solve only one side will lose ground on the other.

> Agents aren't emotional — they need facts. You need to make your entire product presence agent-ready, machine-readable in a way that a logical agent can evaluate and decide whether to surface it to the consumer. It used to be ‘direct to consumer’. Now it's ‘direct to agent’.

Matthias Schlemmer, Partner, Global AI for Consumer Markets Leader, PwC Strategy& Austria

Consumer trust remains the binding constraint on adoption speed. The trust infrastructure is being built now with new protocols for how agents operate emerging; the retailers that build to these standards early will have first-mover advantage when adoption tips.

Retailers are getting strong returns from their AI investments. They see 25% median returns on AI spend — nearly matching cross-sector AI leaders — even while allocating just 7% of functional budgets to AI today.[1] The challenge isn't investment appetite. It's connecting what's already running into end-to-end loops that compound returns across the value chain.

What this means for consumer goods

The traditional growth playbook (brand equity, predictable loyalty, steady margins) is no longer enough. Growth has slowed to single digits year-over-year in many categories. **49% of senior US CPG leaders** doubt their current business structure will survive the decade. Category boundaries are dissolving as consumers seek integrated solutions aligned to lifestyle goals, not products defined by aisle placement. The companies that still rely on controlling shelf space and incremental price increases are losing ground to AI-native disruptors moving with zero legacy drag. And other consumer goods segments like luxury, clothing, and furnishings are also facing significant disruption.

49%

of senior US CPG leaders doubt their current business structure will survive the decade.

Source: PwC's 2025 CPG Executive Survey

1. Sector-level data from **PwC's 2026 AI Performance Study**. Retail and consumer goods findings based on proprietary sector analysis; contact the authors for further detail.

Yet the opportunity for established players is real. Incumbents bring something disruptors cannot easily replicate: deep consumer data, brand trust, and scaled distribution. The question is how to activate those assets in an agentic world.

The data suggests consumer goods companies have further to travel than their retail counterparts. They allocate just 5% of functional budgets to AI (versus 7% in retail and 15% among cross-sector leaders) and see 10% median returns — placing them in the bottom quintile of all sectors for overall AI fitness.[2] The gap is driven by particularly weak innovation capabilities and limited cross-sector collaboration. Only 40% of consumer goods respondents say AI innovation is supported by designated owners embedded in business units, versus 47% in retail and 62% among AI leaders.[3] Without business ownership, pilots stay in the lab. With it, they connect to the enterprise value loops that drive concrete value creation.

[2] Sector-level data from PwC's 2026 AI Performance Study. Retail and consumer goods findings based on proprietary sector analysis; contact the authors for further detail.

[3] Ibid.

02

Value comes from enterprise value loops

Enterprise value from agentic AI comes from taking a process-first view of the business and identifying the end-to-end flows that matter most to revenue, cost, speed, and service. In retail and consumer goods, those flows cut across functions, from demand creation and product launch to pricing, planning, fulfilment, service recovery, and channel execution. This is where agentic AI creates the most value: not in isolated use cases, but in reshaping the core business processes that drive performance.

That distinction matters because siloed use cases rarely change economics at scale. A content agent may improve campaign throughput. A pricing agent may improve decision speed. A demand-sensing agent may sharpen forecasts. But the larger gains appear when those capabilities are connected across a wider value loop, so that signals, decisions, and actions move together rather than stalling at functional boundaries. In this type of broader model, AI does more than automate tasks. It rewires how the business operates.

Enterprise value loops are the right unit of analysis and redesign, because they show where disconnected activities need to become coordinated workflows, where data and judgment need to move faster, and where reusable capabilities can support transformation across the organisation. The firms that capture the greatest value will be those that redesign the few value loops that matter most, then build the foundations to replicate that success across others.

In this section, we review eight value loops that are central to cutting costs and increasing revenues.

AI “Value loops”

Illustrative examples for retail and consumer goods companies

Key:
Cost
Revenue & Cost

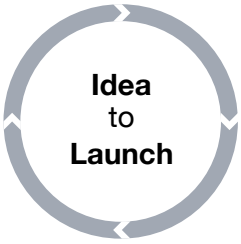

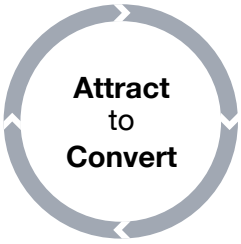

The Rationale

Do next

A key growth driver, as faster, higher-quality product launches reduce failure rates and improve speed to market

Do now

A direct growth driver, as linking targeting, offers and content to sales improves conversion while reducing wasted spend

Do next

A major profit opportunity, as closed-loop optimisation improves pricing, promotions and markdown effectiveness within defined guardrails

Do next

Another major profit opportunity, as better stock accuracy reduces lost sales and cancellations while lowering fulfilment cost

The Rationale

Do next

A major productivity lever, as coordinated execution across stores, digital and fulfilment reduces inefficiency and improves consistency

Do now

The largest controllable cost driver, as store execution directly drives availability and reduces downstream service demand

Do now

A bankable back-office opportunity, as invoice processing is highly repeatable with clean unit economics and clear cost reduction potential

Do now

A major cost-to-serve driver, as high volumes of customer contact and repeat chasing drive rework, refunds and erosion of trust

Idea to Launch

This loop covers the full product lifecycle from design to development, sourcing, and go-to-market. It is one of the clearest tests of whether a company is using AI to speed up one task or redesign a whole process.

Faster launch cycles matter in both sectors. Retailers can offer smaller and sharper ranges, more frequently — differentiating particularly in fashion, beauty and grocery. Consumer goods companies need to move from concept to shelf without the delays created by functional handoffs between R&D, marketing, and supply chain. AI powers multimodal trend sensing, generative concept development, intelligent supplier and material selection, and rapid commercial modelling. The loop closes with the post-launch performance review feeding back into the next range cycle.

> If you can break down the organisational silos from concept all the way to customer, that drives a step change in revenue growth. Concept to shelf can take over 6 months in retail. Agentic orchestration can compress that timeline dramatically, but only when you treat it as one concept to customer journey, and not a collection of isolated automations.

Anna Bancroft, Partner, Consumer Markets & CTO Service Excellence, PwC UK

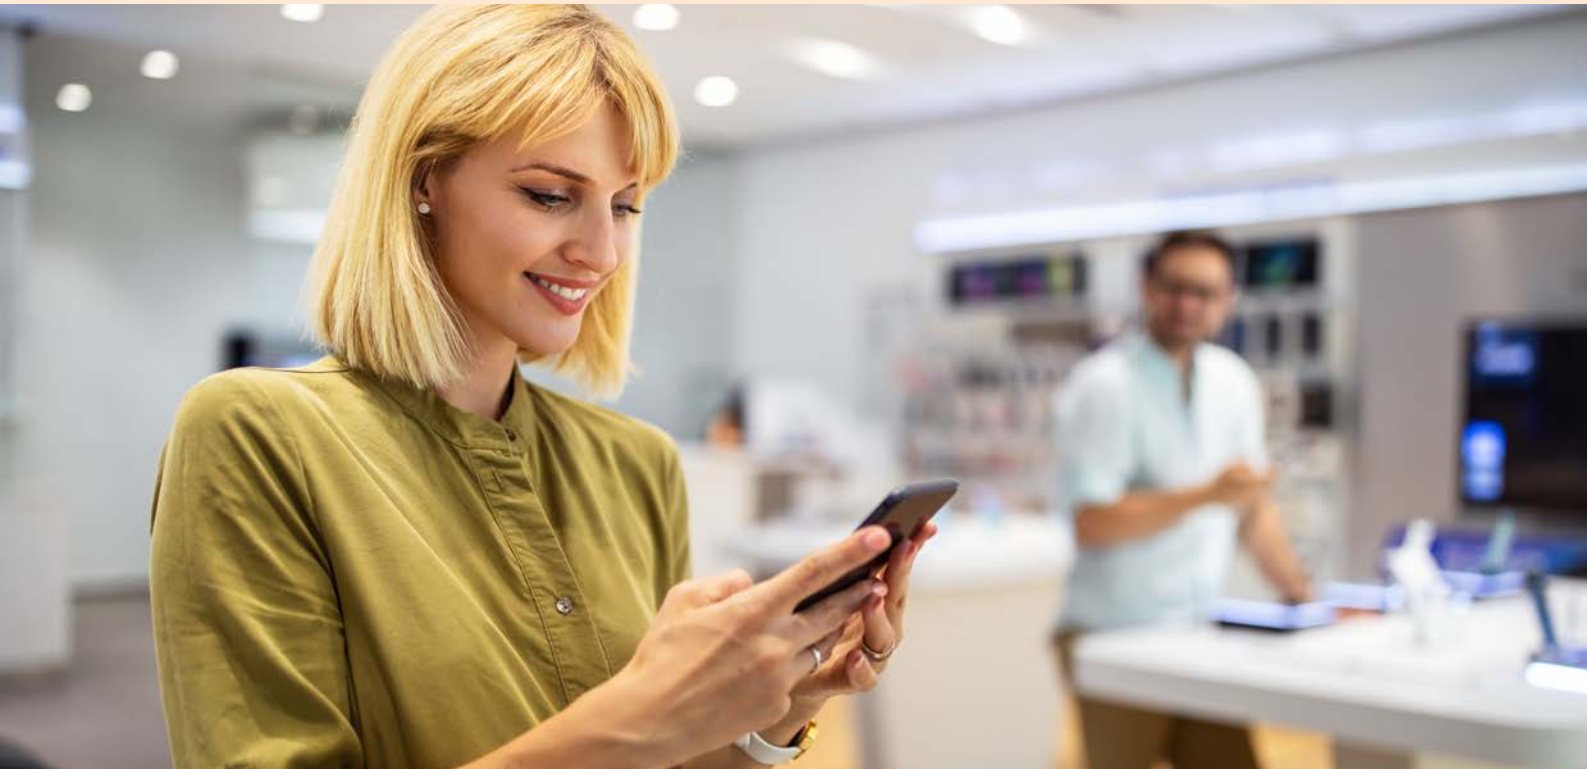

Attract to convert

This loop links content, media, merchandising, digital commerce, and conversion. The cycle covers audience segmentation to customer targeting, content creation, personalised offers and purchase conversion with creative testing that was previously uneconomic.

In an agent-mediated world, discoverability depends less on visual shelf position and more on structured, authoritative product data that machines can parse and trust. Generative engine optimisation is emerging alongside traditional SEO as an increasingly important discoverability discipline. For retailers, that changes the economics of search, recommendation, and conversion. For consumer goods companies, it changes content supply, campaign execution, and how brands show up when an agent, rather than a person, is evaluating the options.

The loop closes when conversion data on which audiences and messages drove sales feeds back to optimise future spend — turning content production and media investment into a learning system rather than a calendar exercise.

Trade to Profit

This is where pricing, promotions, assortment, markdowns, and inventory realities come together. The cycle covers pricing decisions, promotional execution, markdown management and commercial performance review.

Many companies still approach pricing in isolation. The real value comes from closed-loop decisions: adjust pricing within clear guardrails, observe the effect on demand, feed that signal into inventory and replenishment, and adapt again. The loop closes with a strong compounding impact, where each cycle generates insight on price elasticity, cross-product cannibalisation and margin dilution that feeds the next pricing, promotional and inventory decisions. This is particularly important for categories with regular drops and NPD.

In consumer markets, the real opportunity for agentic AI is not just automating individual pricing decisions, but improving decisions and execution across the whole loop — which depends on closer coordination between consumer goods companies and retailers on how value is shared across the chain.

Plan to Promise

This loop connects demand sensing, inventory accuracy, supply planning, and customer promise. It sits at the point where commercial intent meets operational reality. The cycle covers demand forecasting and inventory planning through to availability promises and accuracy of fulfilment commitments.

The value comes from tighter forecast precision (less lost sales from out-of-stocks, less markdown from overbuy), more reliable promises (fewer late deliveries, fewer customer-service contacts) and lower fulfilment cost (fewer split shipments, less expedited freight). For consumer goods companies, the same logic applies upstream: connecting sell-through signals from retail partners to production scheduling and raw material procurement.

Real-time demand sensing grounded in a unified data layer — such as **Microsoft Fabric** — can support more timely demand insights, enabling systems to respond more effectively to changing supply and demand conditions. When POS data, loyalty signals, and warehouse positions sit in one governed environment rather than across fragmented legacy systems, planning teams and their AI tools can work from current signals rather than last week's extract.

Channel to Operate

This loop optimises the role of each channel — stores, online, social, marketplace and wholesale. Routes to market are getting more complex, not less. Retailers now manage stores, e-commerce, marketplaces, social commerce, retail media, partner channels, and increasingly agent-mediated demand. Consumer goods companies face the same fragmentation through direct-to-consumer models, retail partnerships, distributors, and digital commerce platforms. The question is no longer whether a company can optimise each channel in isolation. It is whether it can coordinate execution across channels without adding friction and cost somewhere else in the system.

The agentic value is in the orchestration: aligning labour to footfall, replenishment to sell-through, online stock to store stock, and promotion execution to local conditions. The loop closes with feedback from each channel into the next planning cycle.

Order to Fulfil

This is the execution loop. The cycle covers order placement and fulfilment from picking to packing, dispatch and delivery. Orders still need to be picked, routed, moved, and delivered. Goods still move through physical supply chains. What changes is the cadence and coordination of the decisions around them.

AI enables predictive orchestration of fulfilment nodes, automated status updates and proactive intervention when promises are at risk. Retailers who deliver consistently on delivery expectations are recommended by agents; those who don't get deprioritised. Consumer goods companies face the same logic upstream — suppliers whose fulfilment is predictable and data-rich earn preferred status in automated procurement. Improvements flow downstream with lower customer service requests, lower returns and lower re-delivery cost.

Invoice to Payment

The accounts payable workflow is a strong fit with agentic AI where volume is high, the variance per transaction is low, and exception handling (historically the most expensive part of the cycle) can be managed by specialised worker agents. The highly repeatable transactional cycle covers invoice receipt to three-way matching, exception management, payment execution and payment reconciliation. The loop closes with improvements to the matching rules and the exception classifier.

These may not be the most visible AI stories, but they are often the cleanest economic ones. They reduce cost, improve control, and release funding for harder cross-functional bets elsewhere. In sectors where margin pressure is constant, self-funding is not a side argument — it is often what gives the business room to move on the more ambitious loops.

Issue to Resolution

Customer service is often framed as an efficiency play. That is too narrow. The better lens is service recovery as a loop between issue resolution, trust, repeat purchase, and lifetime value. The cycle covers the full path from customer contact (complaint, request) to resolution (return, refund) and prevention (root cause identification, feedback to functions).

When service agents can see order status, product issues, refund logic, customer history, and likely next best actions in one flow, the business cuts rework and cost while protecting the customer relationship. Strategy& projects that agentic AI could **reduce customer service operating costs by 30%** by 2029. But the bigger opportunity is not cost alone. Customer service pivots from a reactive cost centre into a proactive recovery and retention engine. The loop closes when the patterns of complaints feed back into product, supply chain and operations to prevent the issue recurring. For consumer goods companies, the same feedback loop connects consumer complaints to formulation, packaging, and supplier quality processes — so the system resolves fewer issues over time because it prevents more of them.

Not every loop needs to happen at once. Start with one, prove the economics, then connect it to the next. The compounding happens at the joins: a pricing agent that sees inventory positions, a service insight that reshapes the next product cycle. Many companies stall before they get to that second connection. The next section examines why.

03

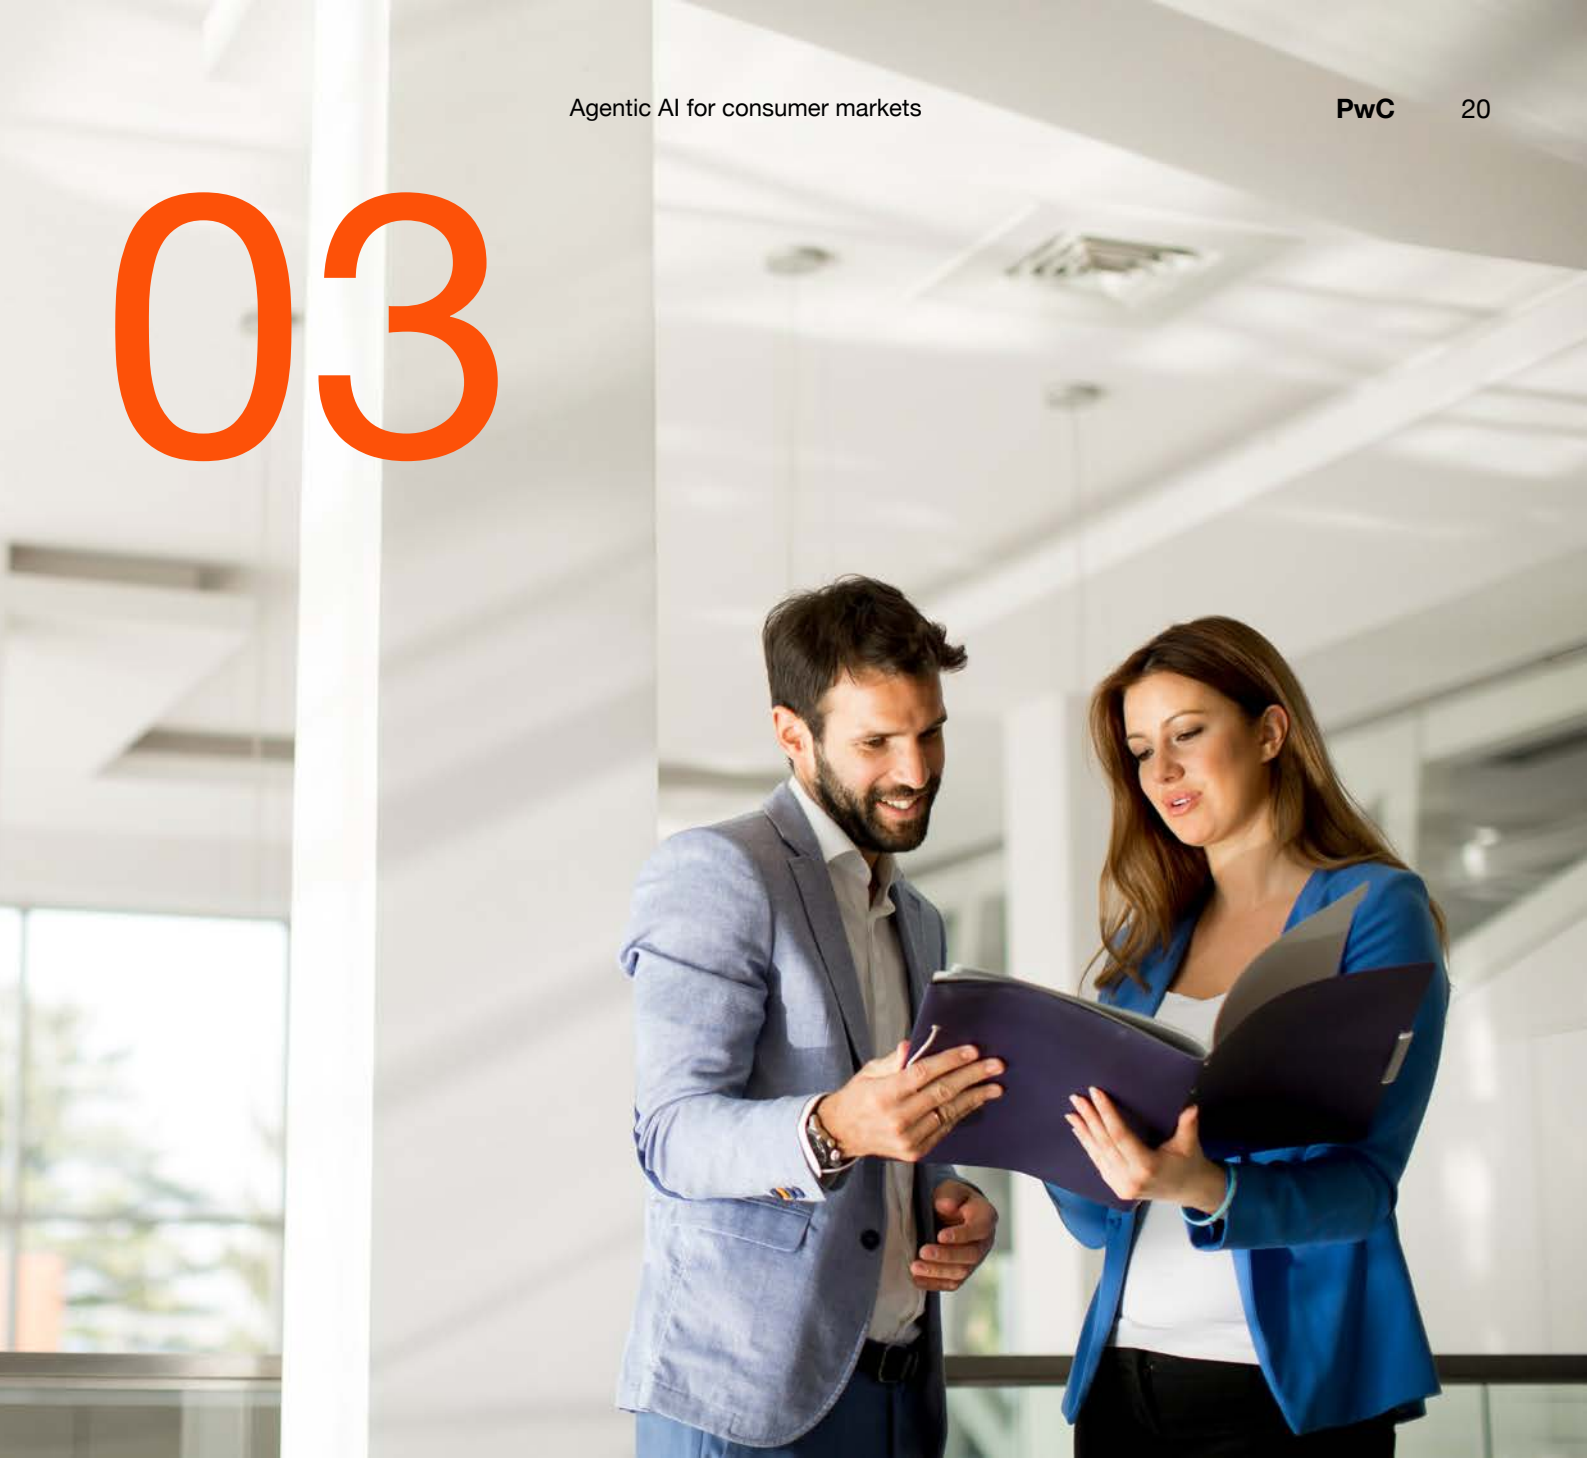

From experimentation to enterprise value: closing the agentic AI gap

of senior executives report broad AI agent deployment

Source: PwC's AI Agent Survey

Most firms now have significant AI activity but little measurable AI impact. Recent **PwC research** shows that more than a third of senior executives report broad AI agent deployment, a further 27% report limited deployment, and 17% say agents are already being used across the full enterprise. Yet another research confirms that few can **demonstrate material revenue impact** from their AI portfolio.

The difference is not who started first. It is who scales.

Agentic AI value shows up at three levels: individual use cases, function-level redesign, and the enterprise value loop. Some organisations work through them in sequence; others skip straight to domain reengineering where the payoff is largest. That last level is the hardest because it demands more than a new model or a new interface. It requires changes to process, data, governance, incentives, and the way people work across boundaries that were never designed to move together.

Unlocking value requires a more holistic approach: end-to-end process redesign, not AI layered on top of workflows that were already broken. Without a North Star strategy, experimentation becomes scattered, producing silos that will never scale. Agentic AI has to become an enterprise decision, not an experimentation phase.

> The biggest value unlocks are no longer about defining use cases, but about tackling points where the organisation feels pain — not where AI feels exciting. Margin erosion, recall risk, complexity, slow speed-to-market: these are places where humans feel overwhelmed and can be greatly assisted by automation, while keeping the human front and centre.
>
> **Nina Lund,** EMEA Consumer Goods & Retail Industry Advisory MD Microsoft

The gap between pilots and performance is persistent because closing it requires fully understanding how agents can drive real business value. It is not mainly a technology gap. It is a redesign gap. When AI is layered onto fragmented processes, the result is faster fragmentation. When it is used to rewire a loop that matters, the economics start to move.

The scale gap

The **most "AI-fit" companies** deliver 7.2x higher AI-driven performance than the rest. Indeed, 20% of firms capture 74% of AI-driven performance. The gap comes from three reinforcing disciplines: they aim AI at growth and reinvention rather than just efficiency, they build focused foundations for priority domains rather than pursuing abstract modernisation agendas, and they embed AI across the enterprise rather than in isolated pockets. When organisations with strong foundations increase AI use, they see nearly double the improvement in AI-driven performance compared to those with weaker foundations.

"AI-fit" companies deliver

7.2

times higher AI-driven performance than the rest.

Source: PwC's new AI performance study

Among cross-sector AI leaders, median AI spending is 5% of revenue. In both retail and consumer goods, it's 1%.[4] But as described above, their readiness profiles differ sharply: retail is investing at leader levels within its functions and seeing returns, while consumer goods lags on both dimensions. The pattern across sectors is consistent: AI leaders spend more, with more discipline. For consumer goods executives in particular, the question is how to close a gap that widens each quarter.

4. Sector-level data from **PwC's 2026 AI Performance Study**. Retail and consumer goods findings based on proprietary sector analysis; contact the authors for further detail.

Microsoft describes the **organisations that are pulling ahead** as "Frontier Firms" — businesses operating at the leading edge of AI adoption, where human expertise and agentic AI collaborate to deliver on-demand intelligence at scale. In consumer markets, Frontier Firms don't just have more agents. They have better-connected agents: ones grounded in unified data, governed by clear rules, and deployed across loops that drive P&L impact.

These firms also integrate principles of responsible AI throughout all aspects of agent development, rather than seeing it as a separate workstream. Strong governance allows agentic systems to scale safely. When agents act with autonomy on behalf of the enterprise and its customers, trust is not a compliance overhead; it is the license to scale. That means embedding transparency, accountability, and human oversight into agent architectures, defining clear escalation paths, making agent decisions traceable, setting explicit boundaries for autonomous action, and continuously monitoring agent behaviour against business and ethical expectations.
The organisations that treat responsible AI as a design principle, not an afterthought, are the ones able to expand agent autonomy with confidence as reliability is proven.

04

Scaling AI means redesigning how work gets done

Most barriers to scaling AI sit inside the organisation rather than in the technology itself. The real test is whether companies can redesign roles, incentives, and management practices fast enough to absorb what AI makes possible.

PwC's 2026 Global AI Jobs Barometer shows that skills sought by employers are changing more than twice as fast in the most AI-exposed occupations — a gap that has widened since last year's Barometer. Since 2022, the most AI-exposed companies have tripled their lead in workforce productivity growth over the least exposed, and workers with AI skills now command a 62% wage premium (up from 57% (a year earlier). In consumer markets, the picture is sharper still: **workers with AI skills command a 118% wage premium**, the highest across all sectors, suggesting that AI capability remains especially scarce and valuable in the sector.

118%

wage premium for workers with AI skills

Source: PwC's 20026 Global AI Jobs Barometer

In consumer markets, the shift is from execution to judgment. Store and field roles increasingly revolve around handling exceptions and interpreting AI outputs. Category and trade teams move from periodic planning to continuous oversight of AI-supported decisions.

> If you truly rethink the process, it changes how people work fundamentally. It's not just putting an email through a company chatbot. You control the agent, evaluate its output, handle what it can't. Design it right, and the old process simply doesn't work anymore — so people have to use the new one.

Matthias Schlemmer, Partner, Global AI for Consumer Markets Leader, PwC Strategy&, Austria

The role that emerges is less "user of a tool" and more what **Microsoft now calls an Agent Boss:** “someone who builds, delegates to, and manages agents to amplify their impact”, so they can work smarter and scale faster. Employees are not the main source of resistance. PwC’s **2025 Global Workforce Hopes and Fears Survey** found that 54% of workers have used AI in their jobs in the past year, and most of those users say it improves productivity and work quality. The bigger issue is whether organisations create the conditions for that usage to become routine.

The Microsoft 2026 Work Trend Index confirms this. It reports that organisational factors like culture, manager support, and talent practices account for more than twice the self-reported AI impact of individual mindset and behaviour. When managers actively model AI use, their teams report significantly higher trust in agentic tools — a bigger effect than any training programme or tool rollout. But behaviour change won't happen while incentive schemes reward the old way of working. **Only 13% of AI users** say they're rewarded for the reinvention of work with AI, even if results aren’t met.

This has direct implications for both retailers and consumer goods companies. A store associate won't use an AI clienteling tool designed to help them build lasting relationships with customers if they are still measured purely on transactions per hour. A trade marketing analyst won't redesign a workflow around agentic sensing if incentives still reward the old cadence.

And PwC's research shows the gap may continue to widen: in 2024, **only 2.1%** of job postings in consumer markets explicitly required AI skills.

The workforce challenge in scaling AI is less about headcount and more about putting the right systems in place. Three actions can help drive success:

- Redesign metrics and incentives around new ways of working; people won't adopt AI at scale if they are still measured on the old process.
- Equip managers to model the behaviour, because adoption follows local leadership more than central messaging.
- Focus skills investment on the value pools where human-agent collaboration has the greatest commercial leverage.

Training matters, but it is not enough on its own. Realising the value of AI requires redesigning the operating model — and that only works when organisation, process, governance, people, and technology move together.

05

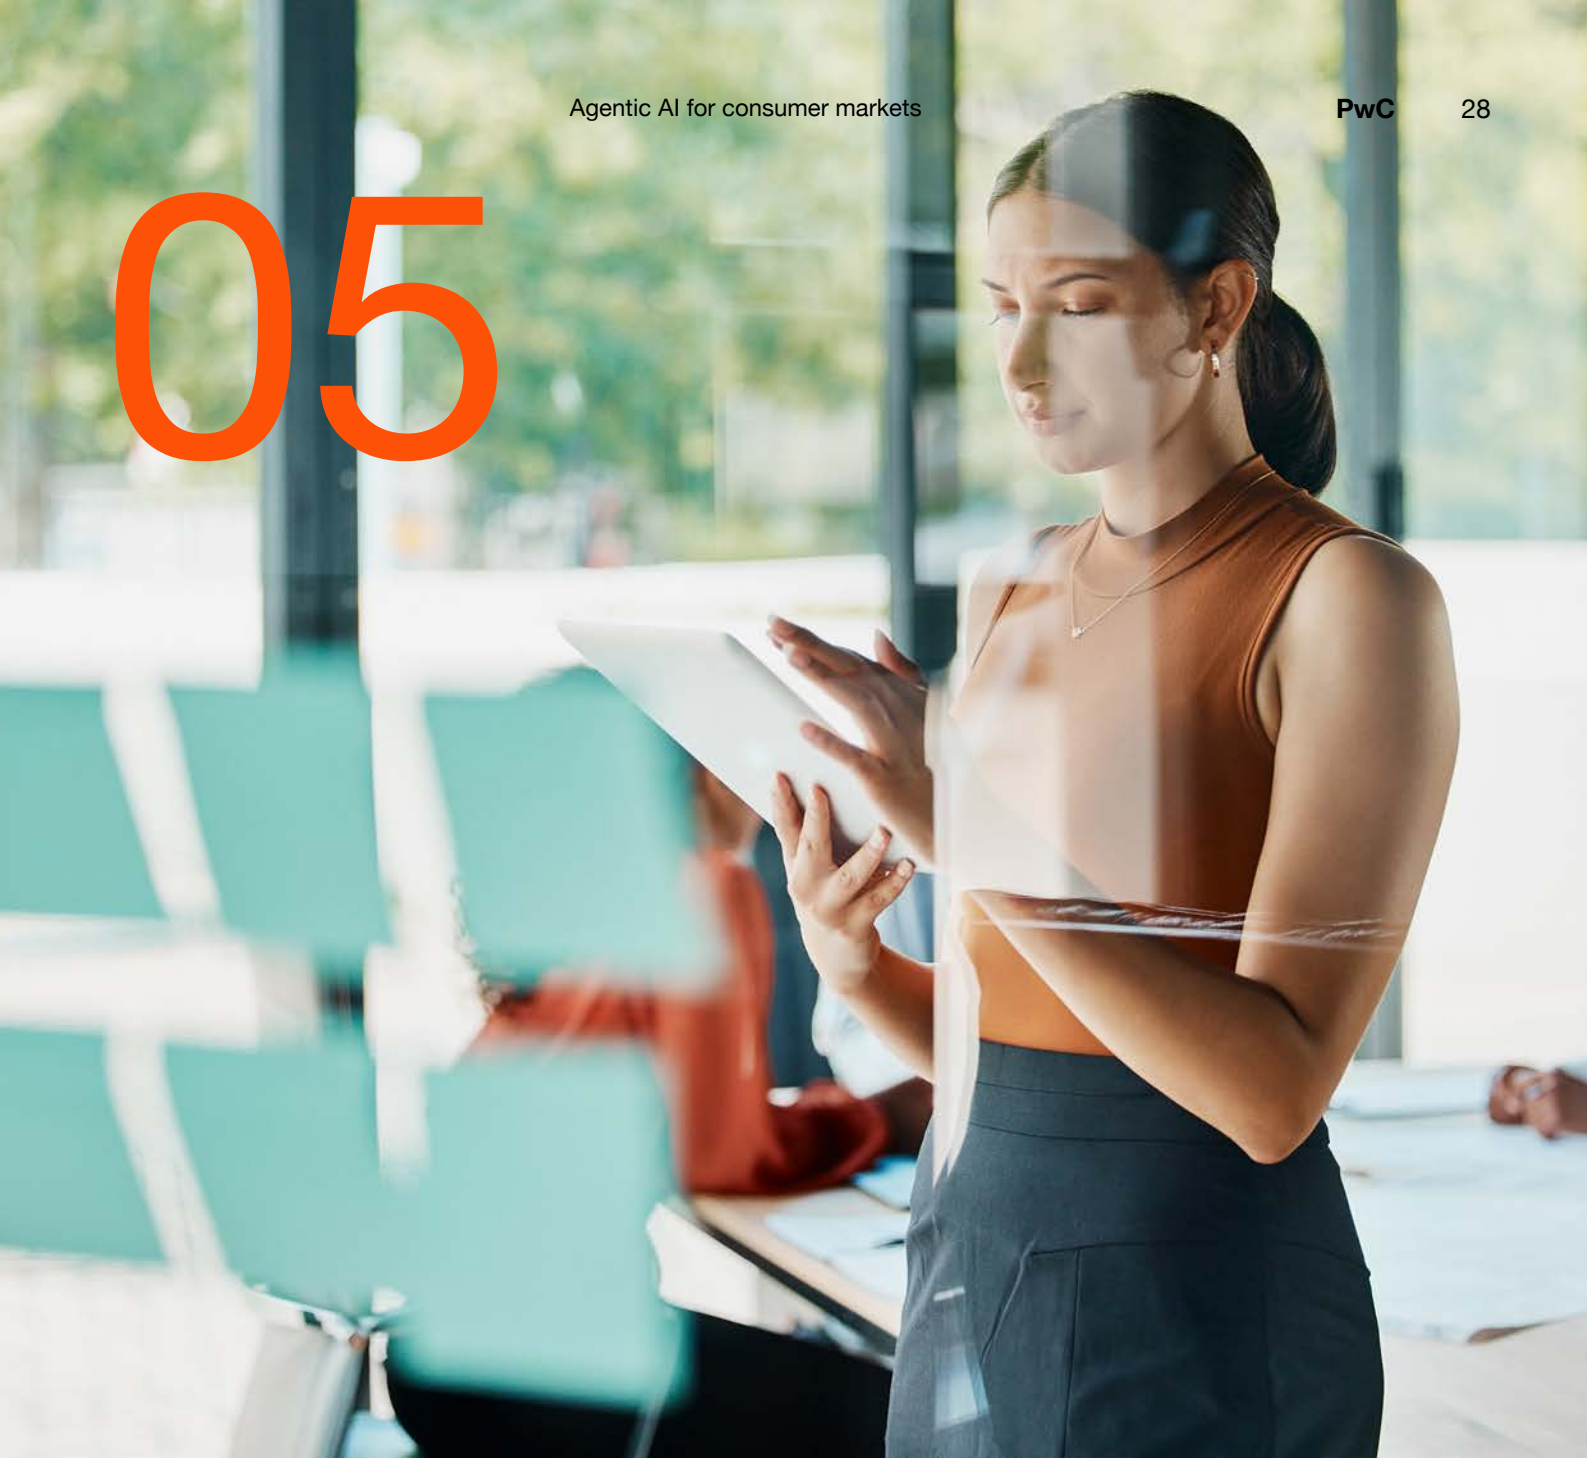

An agent-ready enterprise has three connected layers

The previous sections describe where value sits, what separates leaders from the rest, and why workforce is the binding constraint. But what does it look like in practice when all of these elements come together? How does a retailer or consumer goods company actually operate once it has moved beyond pilots?

The answer is an architecture for an agent-ready operating model: a set of connected layers where agents, data, and people work together across the value chain. Most organisations today have fragments of this. An agent handling customer queries here, a demand-sensing model running there, a pricing tool operating in isolation. The shift from pilots to performance happens when these fragments connect into a coherent operating model.

That operating model has three functional layers, each connected by the human judgment and oversight that runs through all of them, and underpinned by a strong data foundation. These execution layers are enabled by six foundational bedrocks — data, technology, governance, organisation, process and people — which together make the operating model agent-ready.

These three layers align with the Microsoft enterprise AI architecture — Work IQ , Fabric IQ, Foundry IQ and Web IQ — which together form what Microsoft calls **“Microsoft IQ”**. We map
the specific alignment below.

Layer 1:

Experience – where agents meet people

This is where demand forms and where people interact with agents – whether they are customers, colleagues, or supply chain partners. Externally, AI agents sit between customer intent and commercial execution, synthesising options, routing demand, and executing transactions. For example, an agent can ask customers for input on their specific needs and preferences and suggest appropriate products. Brands compete not on visual shelf position but on data quality, structured trust signals, and fulfilment reliability. Internally, agents are the interface through which employees access enterprise knowledge, trigger workflows, and collaborate with AI on daily tasks – from a store associate checking stock allocation to a buyer asking an agent to summarise supplier performance.

The experience layer isn't only customer-facing; it feeds real-time demand signals – commercial and operational – back into the enterprise. When a synthetic signal identifies rising interest in a specific product attribute, or when a colleague agent detects an exception pattern across multiple orders, the deeper layers respond dynamically.

The workforce intelligence layer from **Microsoft IQ** gives agents access to organisational context — emails, meetings, files, and collaboration patterns — so they can operate with awareness of how work actually happens. That makes it possible for people and agents to collaborate across daily operations, whether the “person” is a shopper, a frontline colleague, or a supplier.

Layer 2:

Orchestration and enterprise intelligence

No single agent operates in isolation. The orchestration layer translates intent into multi-step workflows, coordinating and delegating tasks across worker agents while tracking end-to-end progress. A demand-sensing agent talks to a pricing agent, which talks to a supply-planning agent. The orchestration layer manages handoffs, resolves conflicts, and maintains the governing logic that determines what each agent can and can't do. Humans remain in the loop here — setting rules, handling exceptions, and adjusting the governing logic as the business evolves.

This is also where governance scales. As agent ecosystems grow, a control plane monitors agent activities, enforces guardrails, and maintains auditability across the enterprise. Tools such as Microsoft Copilot Studio and Microsoft Foundry accelerate the build; Microsoft Agent 365 acts as the system of record and control for agent governance at scale.

This maps to what Microsoft calls **Foundry IQ**: the layer that grounds, governs, and scales agents securely in production.

Layer 3:

Enterprise systems and operations

This layer is where specialised worker agents execute discrete, repeatable tasks by interacting with enterprise systems via domain APIs — gathering information, validating data, and completing transactions. Those systems (ERP, commerce platforms, supply chain management, CRM, HR) and the physical operations they coordinate (fulfilment, manufacturing, logistics, trade execution) form the foundation these agents act upon.

The fundamental physics don't change; goods still move through supply chains. But the intelligence coordinating those operations runs at a different cadence. Real-time signals replace batch planning. AI and machine learning embedded within core systems apply intelligence at each stage of the end-to-end workflow, using trusted enterprise data.

This maps to what Microsoft calls **Fabric IQ**: connected, contextualised data that grounds AI-driven decisions in a unified, governed estate. In practice, that can mean POS data, customer records, supply chain positions, and pricing rules sitting in a single environment rather than scattered across the legacy systems many retailers still run. Microsoft Fabric can serve this role — bringing fragmented data estates into a governed layer on Azure without requiring wholesale replacement of existing systems.

Agents operate in every layer — their type changes, not their presence

Experience agents interact with people. Orchestration agents coordinate multi-step workflows. Worker agents translate goals into system-level actions. What shifts across layers is scope of autonomy and proximity to human interaction, not whether agents are present.

Human judgment runs across all three layers

No layer operates without people. The front-end needs employees who can design agent experiences and interpret what demand signals mean for strategy. The orchestration layer needs people who can set the rules, resolve exceptions, and adjust the governing logic as the business evolves. The execution layer needs operators who can work alongside AI-coordinated systems and exercise judgment when the unexpected happens.

Each layer generates new roles, new skills requirements, and new ways of measuring performance. The organisations that map their workforce strategy to these layers — as a first-order design decision — are the ones that scale.

> Frontier Firms start with their human capital. Before anything else, AI transformation is an opportunity for people to learn, experiment, and add the tacit knowledge they have of their internal processes.

Nina Lund, EMEA Consumer Goods & Retail Industry Advisory MD Microsoft

The data foundation beneath all three layers

Agents need two kinds of data to function. Internal data (product information, inventory, demand patterns, pricing rules, supplier performance) is what agents reason over and act on. The **end-state is a unified data fabric**: a connected enterprise architecture in which agents can act on complete, cross-channel data in real time. Microsoft Fabric is one platform that can support this architecture, providing a data layer on Azure without requiring replacement of legacy systems.

External data (product attributes structured for machine consumption — pricing, sizing, materials, constraints) is what makes products discoverable to other companies' agents. When a consumer's AI assistant is asked to "order the usual, but less sugar in the soda and sulfate-free shampoo for curly hair," the purchase decision has already been made. Products that can't be parsed won't make it into the consideration set.

The good news: you don't need perfect data everywhere. You need fit-for-purpose, reusable data foundations in the domains where your priority agents operate, then use AI itself to accelerate the data quality work.

06

Five priorities for leaders moving from pilots to performance

> Not starting is almost a statement in itself. Organisations are moving at such a pace that you're missing out on the learnings — the ability to move fast when these things shift from experiment to fast-follow mode.

Anna Bancroft, Partner, Consumer Markets & CTO Service Excellence, PwC UK

The gap between AI leaders and the rest is widening, not narrowing. Companies that move first on reusable foundations, operating-model redesign, and workforce readiness will capture disproportionately more enterprise value.

Making that happen requires disciplined action in five areas:

01

Review the portfolio and classify: scale, redesign, or stop

Take every AI initiative currently running. Scale the ones with strong business cases and measurable results. Redesign those limited to task automation: rethink them end-to-end with multi-agent orchestration in mind. Stop those with no business owner or no measurable impact after 6+ months and redirect the resources. In PwC's experience, most organisations find that a significant portion of their active portfolio belongs in the stop category. The freed resources more than cover the investment needed to scale the winners.

02

Assign business ownership for priority agent domains

Every agent domain needs a named business owner: someone with P&L responsibility, decision rights, and accountability for outcomes. Not an IT project manager or a Centre of Excellence with no authority. The owner's mandate: define the target operating model for that domain including human-agent teaming, set measurable success metrics, and remove organisational blockers.

This is the single most common gap in organisations stuck between pilot and scale. Without business ownership, pilots stay in the lab. With it, they connect to P&L accountability and scale.

03

Treat data as enterprise infrastructure

Data is the most important foundation for agentic AI. Without the right data, an agent cannot decide anything meaningful. The challenge is two-dimensional: internal enterprise data (so agents can operate) and external product data (so agents can discover and recommend).

Internally, organise data into a unified layer for priority domains using platform tooling that abstracts from legacy systems. For many organisations, this starts with modernising core operations on cloud infrastructure — moving from fragmented on-premise estates to a scalable, secure foundation on which AI can operate. Microsoft Fabric provides this capability, bringing fragmented data estates into a single governed environment on Microsoft Azure without requiring replacement of every existing platform. You don't need to solve everything before you start, but treat it as a first-order leadership commitment.

Externally, structure product attributes for machine consumption: pricing, sizing, availability, materials, and constraints in formats agents can reliably parse.

04

Launch one production-grade agent with graduated autonomy

Pick the highest-value domain. Define what "production-grade" means: governance guardrails, monitoring, escalation paths, defined human-agent handoffs, and measurable KPIs (revenue impact, cost reduction, cycle time, customer satisfaction). Each production agent should emit operational telemetry (latency, decision accuracy, exception rates, cost per action) that feeds back into governance and continuous improvement.

Tools such as Microsoft Copilot Studio, with its library of prebuilt agents and built-in security and governance controls, can accelerate the build.

The autonomy question requires mapping each agent's scope against two axes: the consequence of error, and the reversibility of the action. Dynamic repricing of a mid-range SKU based on competitor signals? High-frequency, low-consequence, easily reversed: full agent autonomy is appropriate. Approving a multi-million-euro supplier contract? Low-frequency, high-consequence, hard to reverse: keep a human in the loop. Trust builds incrementally, and so should autonomy. Start with lower-autonomy decisions and expand as reliability is proven.

05

Redesign one team's operating model around human-agent collaboration

Pick a team in a high-value pool: a category management team, a customer service unit, a trade marketing squad. Don't simply provide AI tools; redesign workflows, how performance is measured, what decisions the agent handles versus what humans handle, and how insights loop back. Use this team as the blueprint. Capture what works and codify it. Scale the pattern to the next team, and the next.

The trajectory is clear: from AI assistants to digital labour embedded in every business process. Frontier Firms are already operating this way. The firms that join them will be the ones that converted two or three agents from prototype to production, built the data foundation and workforce model to support them, and created the conditions for it to happen again, faster, every quarter. Each production agent generates data and feedback that makes the next deployment faster, more accurate, and cheaper to operate.

In retail and consumer goods, where margins are thin and competition is fierce, that compounding effect is what will separate the companies that scale AI from those that fall behind.

07

How PwC can help you scale with Microsoft technology

For organisations ready to move from pilots to scaled adoption, the question is no longer where agentic AI might help. It's where to focus first, what foundations need to be in place, and how to turn early momentum into business value. PwC's global Alliance with Microsoft brings those elements together, combining deep sector and functional expertise with the technology, data, security, and AI capabilities needed to scale with confidence.

> AI is not a technology topic and it's not a pure business topic — it's both. Every client conversation starts the same way: 'We have dozens of pilots — which ones go to production?' The answer is never just about the platform or just about the operating model. It's about connecting both.
>
> **Mauro Xavier,** Partner, EMEA Microsoft Alliance Leader, PwC Spain

In consumer markets, that means looking beyond isolated applications and focusing instead on the value loops that drive performance across the business, from customer engagement and commercial execution to supply chain, finance, and workforce productivity.

The AI Agents Accelerator

From prototype to production: scaling agentic AI with confidence

Many organisations have identified where AI agents can create value, yet find it challenging to progress from experimentation to governed, enterprise-grade deployment at scale.

PwC's AI Agents Accelerator helps bridge that gap. A coordinated team of AI agents evaluates and prioritises opportunities across value, feasibility, and risk—underpinned by PwC's industry experience and proprietary process taxonomy, with tailored insight for retail and consumer goods.

The outcome is an implementation-ready blueprint, built on the Microsoft AI stack, with architecture, governance, and integration considered from the outset—enabling leaders to move with pace and confidence from validated opportunity to deployment at scale.

PwC brings deep knowledge of retail and consumer goods value chains, customer expectations, and operating models.
We help organisations identify where agentic AI can strengthen end-to-end performance, translate that into clear business cases, redesign workflows, operating models and governance, and prepare people for new ways of working.

Microsoft provides the enterprise-grade foundation needed to build, deploy, govern, and monitor agents securely and at scale. That includes Microsoft Fabric for the unified data layer, Copilot Studio for building and orchestrating agents with built-in governance, Agent 365 for enterprise-wide control and auditability, and the Microsoft IQ architecture that connects how people work, how data flows, and how agents are governed into a single intelligence layer. PwC is one of the largest global adopters of Microsoft technologies — with Microsoft 365 Copilot deployed across more than 100 countries — and we apply that first-hand implementation experience directly to our client work.

PwC leverages its global alliance with Microsoft to combine deep sector expertise with Microsoft's technology, data, security, and AI capabilities. PwC supports organisations in choosing the right priorities, building reliable foundations, and moving from experimentation to adoption across the value chain. The goal isn't more pilots—it's measurable business impact.

100+

countries

PwC is one of the largest global adopters of Microsoft technologies, with Microsoft 365 Copilot deployed globally

Source: Case study: PwC scales Microsoft Copilot for enterprise AI

Ready to move from pilots to performance? Contact any of our authors to explore how PwC can support your agentic AI journey, leveraging Microsoft technology.

Authors

PwC

Matthias Schlemmer

Partner, Global AI
for Consumer Markets Leader

PwC Strategy&, Austria

Anna Bancroft,

Partner, Consumer Markets & CTO
Service Excellence

PwC UK

Mauro Xavier,

Partner, EMEA Microsoft Alliance
Leader

PwC Spain

With contributions from

Nina Lund,

EMEA Consumer Goods &
Retail Industry Advisory MD

Microsoft