

The Machine Customer

Is UK retail ready for agentic shopping? Context for boards. An operating system for teams. Evidence and action for individual retailers.

A new customer is entering retail

It may never visit a store, scroll a product page or see the campaign created for it.

AI shopping agents are beginning to perform parts of the customer's commercial job: discovering options, comparing propositions, interpreting restrictions, preparing baskets, acting within a mandate and resolving what happens next.

Retailers have spent decades optimising for people and years optimising for search and recommendation algorithms. Agentic commerce introduces a third commercial surface with different requirements. The agent needs exact identity, current offer truth, applicable policy, explicit authority and a route to recourse.

THE UNCOMFORTABLE QUESTION

If a machine customer tried to shop your business today, would it reach the right proposition, understand it correctly, act within the customer's authority and resolve the consequences?

Why this belongs on the board agenda

Revenue: delegated demand may be won or lost before a human reaches the retailer.

Control: the retailer can appear in agent answers even when it blocks the agent's preferred route.

Trust: wrong products, stale offers or missing conditions can travel at machine speed.

Accountability: access, data, trading, security, legal, service and commerce all own part of the outcome.

The agentic shelf is not another website channel

It is the connected evidence and authority that allow a machine to represent a customer safely.

Three shelves, three different design problems

Source: The Future Collective framework

Human surfaces support judgement and persuasion. Algorithmic surfaces shape ranking and recommendation. Agentic surfaces must preserve identity, evidence, offer truth, policy, authority, transaction and recourse across a delegated outcome.

STRATEGIC IMPLICATION

A retailer can perform well on the human and algorithmic shelves while remaining unreadable, ambiguous or unactionable to a machine customer.

Readiness must survive the complete customer job

A successful page load is not the same as a successful delegated outcome.

Find, Understand, Act and Resolve

Source: The Future Collective framework

What each gate asks

GATE	EXECUTIVE TEST
Find	Can the agent reach and identify the exact proposition through an authorised route?
Understand	Can it interpret product, offer, eligibility, policy, evidence and uncertainty correctly?
Act	Can it hand off or transact within a specific, evidenced customer mandate?
Resolve	Can the customer change, cancel, return, dispute or escalate with continuity of context?

THE PROPAGATION PROBLEM

A failure at an early gate travels downstream. A perfect checkout interface cannot repair the wrong proposition; a safe refusal does not prove the exact retailer source was found.

Blocking agents does not remove the retailer from agentic commerce

It may remove the most authoritative source while leaving third-party and stale material available.

Six failure modes boards should expect

Illustrative exposure map

FAILURE MODE	BUSINESS CONSEQUENCE
The authorised agent is blocked or challenged	The retailer becomes harder to verify than an intermediary or scraper
Near-identical products cannot be distinguished	The wrong pack, variant, seller or fulfiller is recommended
Price or promotion conditions are incomplete	A headline price becomes commercially misleading
Availability has no location or timestamp	The agent converts uncertainty into a false promise
Policy is prose without product linkage	The wrong restriction or right is applied to the proposition
The action cannot prove mandate or recourse	The retailer cannot reconstruct, contain or remedy the outcome

What not to conclude

Blocked is not absent. Readable is not legible. A recommendation is not an authorised action. A safe outcome is not proof of source fidelity. The research and diagnostic method keeps observation, inference and commercial implication separate.

THE ISSUE IS ARCHITECTURAL

This is not an SEO problem, a bot-policy problem or a chatbot problem. Each may be involved, but none owns the full customer outcome.

Readiness is a ladder, not a yes-or-no label

The goal is to make delegated outcomes progressively more observable, usable and trustworthy.

Six levels of agentic-commerce maturity

Source: The Future Collective framework

Maturity is evidenced through representative customer missions. A deployed API, published policy or new platform counts only when the intended mission behaves as specified and the decision can be reconstructed.

SCORING DISCIPLINE

Control maturity, mission outcome and evidence confidence are recorded separately. One optimistic average cannot conceal a failed high-consequence gate.

The operating model crosses the organisation

Readiness emerges when six capabilities tell the same commercial story.

The six workstreams

CAPABILITY	OUTCOME
1. Access and channel policy	Authorised machine customers can reach the right surfaces under an explicit, monitored policy
2. Canonical product identity	The exact proposition, variant, offer, seller, fulfiller and document resolve to one object
3. Offer truth and parity	Price, conditions, stock, fees and fulfilment remain coherent through commitment
4. Policy, safety and recourse	Applicable rules, stop conditions, rights and escalation are governed decision inputs
5. Transaction and authority	The retailer distinguishes agent, customer and mandate, and can audit or recover actions
6. Governance and measurement	Ownership, evidence, incidents, investment and continuous testing operate together

What the licensed playbook adds

Each workstream is translated into auditable controls with a primary owner, implementation sequence, minimum evidence pack, definition of done, assurance test, dashboard measures and first 30-day move.

SAMPLER BOUNDARY

The full 30-control catalogue, acceptance criteria and 14 editable templates are deliberately reserved for licensed buyers.

Test real delegated missions, not generic prompts

A defensible answer fixes the proposition, context, authority and success criteria before the test begins.

A representative mission specifies

The customer's delegated job

The exact proposition or selection constraints

Location, time, account state, client and route

The authority boundary and confirmation events

Expected outcomes at Find, Understand, Act and Resolve

Stop conditions, safe unknown behaviour and evidence to retain

Every observation receives an evidence state

Evidence states

STATE	MEANING
Confirmed present / absent	Directly established within the defined method and context
Blocked or challenged	Access was prevented or diverted; absence remains unproven
Sign-in required / region unavailable	A contextual boundary prevented the intended test
Unknown / not applicable	Evidence is insufficient, conflicting or the criterion does not apply

WHY IT MATTERS

A retailer can safely use the result in a board decision because the evidence, context, limitation and reviewer travel with the conclusion.

Create a funding decision, not a transformation theatre

The objective is to establish current truth, contain exposure and sequence the right controls.

Four-phase mobilisation

PERIOD	EXECUTIVE QUESTION	EXIT EVIDENCE
Days 0-30	What are we permitting, testing and who owns the outcome?	Mandate, scope, mission portfolio and governance
Days 31-60	What does the machine shelf actually do?	Evidence ledger, gate results and initial maturity
Days 61-75	Why does it fail and what is exposed?	Root causes, exposure map and explicit unknowns
Days 76-90	What should we fund, stop or pilot?	Stabilisation plan, roadmap and board decisions

The retailer leaves with

- A machine-channel map
- An evidence-backed maturity profile
- A failure and exposure map
- A root-cause map
- A 90-day control plan
- A 12-month investment roadmap

PRIORITY RULE

Fix a material failed gate before polishing a later one. Fund outcomes and assurance tests, not initiatives described only as AI capability.

Buy the level of answer the organisation needs

Start with shared context, move to implementation, then establish retailer-specific evidence.

Commercial routes

ROUTE	WHAT IT PROVIDES
The Machine Customer - flagship report	UK context, evidence, board implications and the category framework, as a designed edition.
Ready for the Machine Customer - implementation playbook	Thirty controls, fourteen working templates and the 90-day mobilisation, sequenced for cross-functional teams.
The engagement - bespoke, per retailer	Retailer-specific evidence, exposure and the funded roadmap, generated from recorded probe data and scoped per client.

The product stack

The report answers why. The playbook answers how. The engagement answers where your retailer fails, what it should fund, and whether it is improving.

ENTERPRISE LICENCE

All materials are licensed for defined internal organisational use. The playbook and templates are not supplied as unrestricted downloadable content.

One executive can sponsor the decision; no single function can implement it

The strongest first conversation includes the leaders who own customer outcome, commercial truth and operating control.

Who should be in the room

ROLE	QUESTION THEY OWN
Executive sponsor	Is this strategically material, and who can resolve cross-functional trade-offs?
Digital / ecommerce	How do machine customers discover and progress through the journey?
Trading / product data	What is the exact proposition and authoritative offer?
Technology / architecture	Which systems, interfaces and identities enable the outcome?
Security	Which machine actors and routes are permitted, challenged or refused?
Legal / compliance / risk	Which claims, policies, authority and safety boundaries apply?
Customer service	How are changes, returns, disputes and escalation resolved?

A useful first meeting should decide

Which customer missions matter most

Which categories create the greatest consequence

Who owns the end-to-end outcome

Whether the organisation needs context, implementation tools or its own evidence

Start with one uncomfortable mission

Choose a customer job your board assumes the retailer can support. Then ask the machine customer to perform it.

THE QUESTION

Can the agent find the exact proposition, understand the current commercial and policy truth, act within a specific customer mandate, and resolve what happens next?

If the answer is uncertain, the uncertainty is already useful. It identifies an evidence gap, ownership gap or control gap that the retailer can make visible before delegated demand scales.

Three ways to continue

Buy the flagship report to create a shared board and executive context.

Buy the implementation playbook to give cross-functional teams the controls and working tools.

Commission an engagement to establish current truth, exposure and the funded roadmap.

PRIVATE EXECUTIVE BRIEFING

The Future Collective can brief your leadership team on the category shift, test the strategic implications and define the right next level of evidence.

THE FUTURE COLLECTIVE

UK-specific. Independent. Evidence-disciplined. Designed for boards and the teams that must implement the answer.

© 2026 The Future Collective. Executive sampler. Not licensed for commercial reuse or redistribution.