AI USE-CASE REVIEW · ILLUSTRATIVE SAMPLE

Where AI pays at a mid-size food maker, and what has to be true first

A use-case review for a hypothetical US packaged-food company that has AI tools and no shared view of where they pay. Public sources only. Ratings are illustrative and written as words.

Prepared by Eric Palmer Consulting

November 2026

Start with four use cases that are ready now, fix the data under three more, and hold the rest

Of twelve candidate use cases, four can start this quarter on data and owners the company already has. Three pay more but need their data fixed first, three are cheap pilots, and two should wait.

1. **The question has moved from whether to where.** Firms with at least 250 employees use AI at 1.9x the national rate (37% against 19.8%, Census).

2. **Each use case is read four ways.** The value it creates, what has to be true first (the data, the process and an owner), what it costs to keep running, and where a person checks it.

3. **Four are ready now.** Trade promotion reads, consumer email triage, the monthly review draft and a competitor digest. Each has its data, an owner and a named check.

4. **The biggest prizes need their data fixed first.** Demand forecasting and the weekly retailer read pay most, and both depend on data with no single owner today.

5. **Decide three things now.** Approve the first four, name an owner for the forecasting and retailer data, and fund the three pilots on a fixed budget with a stop date.

Sources: S1 (full citations in the appendix).

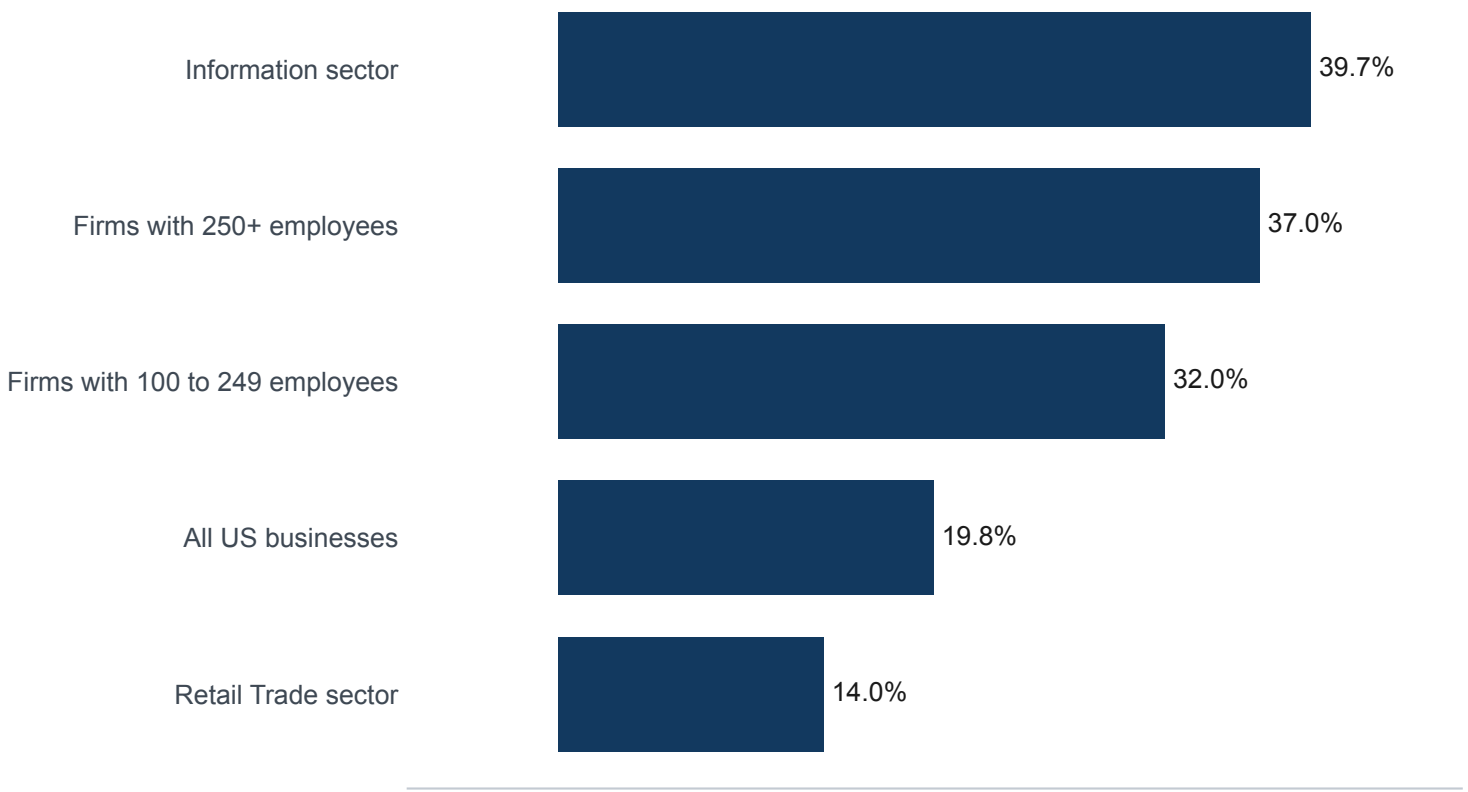
Five questions this review answers, in the order a leadership team asks them

Section	The question	Where the answer is
Why now	Is AI something we should be deciding on this year?	Adoption by firm size
The method	How is each use case judged?	Four lenses
The long list	Which use cases were considered, and how did each rate?	The twelve, rated
The answer	What do we start, fix, pilot and hold?	Four groups, what has to be true, running cost
Next steps	Who does what in the next 90 days?	Options, counter-case, pre-mortem, plan

- Each section ends in a decision. The only market figures are cited, and the books behind the methods are listed at the back.

Firms with 250 or more employees use AI at 1.9x the national rate

Businesses using AI, percent, to May 3, 2026 (Census BTOS)



37%

Firms with at least 250 employees using AI (Census)

19.8%

All US businesses, the national rate (Census)

- More than a third of firms this size already use AI somewhere in the business.
- So the useful question is where AI pays, and what it costs to keep it paying.

Sources: S1 (full citations in the appendix).

The tools are bought, the pilots are scattered, and nobody can say which ones pay

SITUATION

AI tools are licensed and teams are trying them

- Several functions run their own trials
- Licences are paid centrally
- Early users are keen and vocal

COMPLICATION

No shared view of value, readiness or running cost

- Each trial is judged by the team that ran it
- Data problems surface late, after the build
- Nobody owns the cost of keeping a tool useful

RESOLUTION

Rate every candidate the same way and decide in groups

- Four lenses, applied to every use case
- Four groups: do now, fix first, pilot cheaply, not yet
- A 90-day plan with owners

THE METHOD

Every use case is read through four lenses before it is placed in a group

The same questions for every candidate, so a keen team's favourite and a quiet team's idea are judged on the same terms.

VALUE

What it changes, and for whom

- The decision or task it speeds up
- Who notices if it works
- Rated high, medium, low or unclear

WHAT HAS TO BE TRUE

Data, process and an owner

- Is the data in one place and trusted?
- Is the process written down?
- Is one person accountable?

RUNNING COST · B1

Value paired with upkeep

- Licences, review time, data upkeep
- Measured beside the value, never alone

THE CHECK · B2

Where a person signs off

- A named pause point before output is used
- Who checks, and against what

So what. A use case with high value and no owner is not ready. It goes to fix first, with the owner named as the first step.

Frameworks: B1, B2 (books, see appendix).

Twelve candidates, rated the same way, with the group each one lands in

Use case	Function	Value	Readiness	Running cost	Group
Post-event read of trade promotions	Sales and finance	High	Ready	Low	Do now
Consumer email triage and draft replies	Consumer affairs	Medium	Ready	Low	Do now
First draft of the monthly review commentary	Leadership office	Medium	Ready	Low	Do now
Competitor and category news digest	Strategy	Medium	Ready	Low	Do now
Demand forecast support	Supply planning	High	Data in three systems	Medium	Fix first
Weekly retailer sales read	Sales	High	No single data owner	Medium	Fix first
Supplier contract review	Procurement	Medium	Contracts scattered	Low	Fix first
Label and claims copy first drafts	Marketing and regulatory	Medium	Needs legal review	Medium	Pilot cheaply
Consumer research synthesis	Insights	Medium	Ready	Low	Pilot cheaply
Key account call preparation	Sales	Medium	Partly ready	Low	Pilot cheaply
Plant maintenance prediction	Operations	High	Sensors not connected	High	Not yet
Recipe and formulation suggestions	Research and development	Unclear	No test data	High	Not yet

THE FOUR GROUPS

Four start now, three wait on data, three are cheap pilots, and two should wait

DO NOW

Ready, owned and cheap to run

- Post-event read of trade promotions
- Consumer email triage and draft replies
- First draft of the monthly review commentary
- Competitor and category news digest

FIX FIRST

Valuable, blocked by data or ownership

- Demand forecast support
- Weekly retailer sales read
- Supplier contract review

PILOT CHEAPLY · B3

Promising, unproven beyond keen users

- Label and claims copy first drafts
- Consumer research synthesis
- Key account call preparation

NOT YET

Costly to run or no way to test

- Plant maintenance prediction
- Recipe and formulation suggestions

So what. A pilot that works for the keen early users still has to win over the majority who did not ask for it (B3). So each pilot has a fixed budget, a stop date and a test run with ordinary users.

Frameworks: B3 (books, see appendix).

DO NOW

Each of the four has its data, an owner and a named check before anyone relies on it

Use case	The data it uses	Owner	The check before use (B2)
Post-event read of trade promotions	Promotion calendar and shipment data already pulled monthly	Revenue management lead	Analyst signs off the read before it goes to the account team
Consumer email triage and draft replies	The consumer inbox and approved reply library	Consumer affairs manager	Every draft is read and sent by a person. Safety mail goes straight to a person
First draft of the monthly review commentary	The finished monthly numbers	Chief of staff	The function head edits and owns the final text
Competitor and category news digest	Public news, filings and press releases	Strategy lead	Sources linked on every item, read before it circulates

Frameworks: B2 (books, see appendix).

What keeps costing after launch, and the value it is measured against

Each cost is paired with the value it buys (B1), so a use case that saves time but needs more review than it saves shows up as a loss.

Cost after launch	What it is	Paired with
Licences and usage	Seats and usage charges for the tools	Hours saved by the people using it
Review time	The person who checks output before it is used	Errors caught before they reach a customer or retailer
Data upkeep	Keeping the inputs current and in one place	How often the output is trusted without rework
Model and tool changes	Re-testing when a vendor changes its model	Whether quality holds after each change
Ownership	The named person who answers for it	Whether anyone notices when it stops working

Frameworks: B1 (books, see appendix).

OPTIONS

Three ways to run the first wave, and what would have to be true

OPTION 1

Each function runs its own

Functions pick and fund their own use cases with the tools already licensed.

WHAT HAS TO BE TRUE

- Each function has someone who can judge the output
- Data problems stay inside one function
- Costs are tracked function by function

Read: Fast to start, and the scattered-pilot problem continues.

OPTION 2 · RECOMMENDED

A small core, function owners

One person keeps the list, the lenses and the costs. Each use case has an owner in its function.

WHAT HAS TO BE TRUE

- A leader gives the core an hour a week
- Function owners accept the four lenses
- The list is re-rated each quarter

Read: Keeps the judging consistent without a new team.

OPTION 3

A central AI team builds

Hire or assign a team that builds and runs every use case for the business.

WHAT HAS TO BE TRUE

- Enough use cases pay to fund the team
- Functions hand over their processes
- The data fixes are already done

Read: Possibly right later. Premature before the data is fixed.

The strongest case for waiting, and the answer

THE CASE FOR WAITING

- **The tools change every quarter.** Anything set up now may need redoing.
- **Retail is a slow adopter.** Only 14% of retail businesses use AI (Census), so our customers are not pushing yet.
- **The biggest prizes are not ready.** Forecasting and the retailer read both wait on data.

THE ANSWER

- **The first four do not depend on one tool.** They rest on the data, the owner and the check, which carry over when a tool changes.
- **Peers of this size are already moving.** 37% of firms with 250 or more employees use AI (Census).
- **Waiting does not fix the data.** Naming the data owners now is what makes the biggest prizes possible later.

Sources: S1 (full citations in the appendix).

PRE-MORTEM

If the first wave has failed a year from now, these are the likeliest reasons

Written before the first wave starts (B4): assume it failed, write down why, then act on each reason today.

How it failed	The early warning	What we do now
Output was trusted without a check	An error reaches a retailer or a consumer	Every use case ships with its named check
Review took longer than the work it replaced	Review time rises month on month	Review time is measured beside hours saved
The data owners were never named	Forecasting is still in fix first next quarter	Owners named within 30 days, reported to the leadership team
Pilots never ended	A pilot passes its stop date with no decision	Each pilot has a budget and a stop date set today
Only the keen teams used it	Usage stays inside the teams that asked for it	Each pilot is tested with ordinary users before it scales

Frameworks: B4 (books, see appendix).

NEXT 90 DAYS

What happens after the leadership team decides

When	What	Owner	Comes back as
Within 30 days	The four do-now use cases live, each with its check	Each use case owner	A one-page status for each
Within 30 days	Owners named for forecasting and retailer data	Chief operating officer	Names and a data-fix plan
Within 30 days	Three pilots started, each with a budget and stop date	Pilot owners	Pilot scope, for information
Within 90 days	First read of value against running cost	The core owner	The paired measures for each use case
Within 90 days	The long list re-rated	The core owner	An updated list, with what moved and why

SCOPE

What the review delivers, what it needs from you, and what it leaves out

DELIVERED, AND WHAT YOU PROVIDE

- **Decision-ready deck.** The groups, the reasons and the 90-day plan, editable.
- **The rating model.** Every use case and its ratings, so your team can re-rate it each quarter.
- **From you.** The decision it feeds, the processes in scope and one named contact for questions.

OUT OF SCOPE

- **Building the systems.** The review says where to build and what has to be true, not the build itself.
- **Choosing a vendor.** Tools are named only where a use case needs a capability.
- **Legal and security sign-off.** Flagged where a use case needs it, decided by your own teams.

APPENDIX

Sources

ID	Source	What this deck uses	Accessed
S1	U.S. Census Bureau (Adam Grundy, Cory Breaux and Dhanapati Khatiwoda). Large Firms With at Least 20 Employees Biggest AI Users. America Counts story, May 26, 2026. Business Trends and Outlook Survey, December 14, 2025 to May 3, 2026. https://www.census.gov/library/stories/2026/05/ai-use-businesses.html	AI use rates by firm size and sector	2026-11-06

APPENDIX

Frameworks and books cited

Methods named on these pages and where they come from. They shape how the questions are asked. No figure on these pages comes from them.

ID	Book	Chapter	How this deck uses it
B1	Andrew S. Grove. High Output Management. Vintage Books, 2015.	Chapter 2, Managing the Breakfast Factory	Each use case's value is read beside what it costs to keep running
B2	Atul Gawande. The Checklist Manifesto. Metropolitan Books, 2009.	Chapter 6, The Checklist Factory	A named pause point where a person checks the output before it is used
B3	Geoffrey A. Moore. Crossing the Chasm, Third Edition. HarperBusiness, 2014.	Chapter 1, High-Tech Marketing Illusion	Why a pilot that works for keen early users still has to win over everyone else
B4	Daniel Kahneman. Thinking, Fast and Slow. 2011.	Chapter 24, The Engine of Capitalism	The failure story written before the first wave starts

NEXT STEP

Have a decision like this to make?

This sample shows the kind of work I do for consumer brands and their buyers. Send a few lines in writing: the decision, the deadline, and who has to be convinced. Scoping questions come back, usually within one business day, and no call is needed to get started.

Send a written brief

ericwpalmer.com/contact