SMALL BUSINESS OWNERS · PRACTICAL AI WORKSHOP

AI, Analytics & Agents

Learn to tell what an AI tool really is, what it is allowed to do, and who signs off — before it touches a customer, a payment or your books.

CLASSIFY

What is it, really?
Seven roles, four questions.

→

CHECK

What is it running on?
Your numbers, verified.

→

CONTROL

Who signs off?
A named person, every time.

Free Simulation App to Explore and Learn

https://tinyurl.com/95fbejap

Code: SMB-ACTION

WHY THIS ROOM · WHY NOW

Everyone says they use AI. Few can say what kind.

The same small businesses produce very different adoption numbers — because nobody agrees on what counts as AI.

58%

SAY THEY USE GEN AI

of US small businesses in 2025, up from 40% in 2024 and 23% in 2023

19.8%

USE AI IN A FUNCTION

of US businesses in the prior two weeks, Census survey, May 2026

<20%

OF 1–4 PERSON FIRMS

report using AI; firms under 20 staff showed no significant rise

20–23%

PLAN TO START SOON

of businesses expected to be using AI within the next six months

! **WHY THE NUMBERS DISAGREE**

58% counts owners who say they use generative AI — including a free chatbot for one email. 19.8% counts AI used in a business function in the last two weeks. Same economy, different definitions. If you cannot classify a tool, you cannot count it, price it or control it.

YOUR ROLE IN THIS

In a small business you are the buyer, the approver and the person who answers for the output. There is no IT or compliance team between an AI tool and your customer.

US Chamber of Commerce, Empowering Small Business 2025 (n=3,870 firms under 250 staff, self-reported). US Census Bureau, Business Trends and Outlook Survey, Dec 2025–May 2026 (nationally representative).

THE PROMISE

Anyone can prompt. Almost nobody can prove.

This is not a prompting class. It is about what must be true before an AI output changes anything in your business.

WHAT MOST OWNERS DO TODAY

- Ask a chatbot about a tax rule, a wage law or a lease clause.
- Paste a fluent, confident answer into an email, a quote or a social post.
- Buy a tool because the vendor calls it “AI-powered” or “agentic.”
- Have no answer when a customer, accountant or inspector asks where a figure came from.

WHAT YOU WILL BE ABLE TO DO AFTER TODAY

- Classify any AI tool into one of seven roles, whatever the vendor calls it.
- Say what each tool is — and is not — allowed to do on its own.
- Check the numbers a tool runs on before you trust what it builds.
- Put a named person in front of any action that touches money, customers or staff.

The scarce skill for a small business is not building AI. It is knowing exactly what you bought, what it can do without asking, and whose name is on the result.

CASE 1 · WHO OWNS WHAT THE AI SAYS

A chatbot invented a policy. The company owned it.

WHAT HAPPENED

A grieving passenger asked the airline's chatbot about bereavement fares. It told him he could claim the reduced fare retroactively — contradicting the airline's own published policy page. He relied on it. The airline refused to honour it.

THE DEFENCE THAT FAILED

Air Canada argued the chatbot was a separate legal entity responsible for its own actions. The tribunal called that a remarkable submission, and held the airline liable for negligent misrepresentation.

WHY AN $812 AWARD MATTERS TO A SMALL BUSINESS

- The money is trivial. The principle is not: a business owns what its chatbot tells a customer.
- Now transpose it. Your website assistant quotes a refund window, a price, a warranty or a delivery date you never offered.
- Ask the practical question, not the legal one: which AI-written message could a customer reasonably rely on today — and who checked it?

Moffatt v. Air Canada, 2024 BCCRT 149 (14 February 2024) — British Columbia Civil Resolution Tribunal, not a court. Precision matters when you cite it.

CASE 2 · WHEN THE AI ADVISES YOU

A city's AI told business owners to break the law.

Owners often receive a confident AI answer — and acting on a wrong one about wages or tenants is your liability.

Oct 2023

LAUNCHED FOR OWNERS

New York City's MyCity bot was pitched as a guide to city rules for business owners.

Mar 2024

ERRORS REPORTED

Reporters found it saying employers could take a cut of workers' tips and landlords could refuse housing vouchers.

~$600K

REPORTED BUILD COST

It stayed online with a beta warning until the city took it down in early 2026.

THE LESSON FOR AN OWNER

Nothing here needed technical skill to catch. It needed one habit: for any AI answer about wages, tax, tenants, safety or refunds, ask where the rule is written — and read that source before you act.

THE HONEST COUNTERWEIGHT

The city said the bot had given thousands of accurate answers, and it warned that responses could be inaccurate. Both can be true. But a disclaimer does not change the law — it moves the risk to whoever relies on the answer. If you act on it, the risk is yours.

The Markup and THE CITY, “NYC's AI Chatbot Tells Businesses to Break the Law,” 29 March 2024; The Markup, 30 January 2026 (updated 4 February 2026: bot taken down). Build cost as reported, not audited.

TABLE EXERCISE · 8 MINUTES

What failed first?

Pick one primary failure for each case. Defend it with a reason, not a preference.

A — MODEL

It produced an inaccurate or unsupported output.

B — EVIDENCE

It ran on stale, conflicting or unvalidated source data.

C — WORKFLOW

There was no verification or escalation step in the path.

D — GOVERNANCE

A consequential system was deployed without controls.

A STRONG ANSWER

More than one layer failed in both cases — but governance owns the system boundary. The model's behaviour was the symptom. No verified source and no check before an answer reached the public were the mechanism. Putting a confident tool in front of people without controls was the cause.

Each of those four layers maps to a decision a small-business owner can actually make: what information the tool may use, what check comes before it answers, who signs off, and what gets kept on record.

THE VALUE GAP

The bottleneck is not the tool. It is everything around it.

Five simple controls separate an AI tool that impresses in a demo from one you can safely let near customers and money.

1

One agreed definition for each key number

Written down, and the same in your POS, your books and the report you read.

2

A checked source of truth

Counts and totals compared with your bank, POS or books before anything is built on top.

3

A check before anything goes out

A defined check that must pass — not a hope that someone will glance at it.

4

A named person who signs off

You or a specific employee — not “the team” — accountable for any action that matters.

5

A record of what happened

What went in, what came out, who approved it, and when. Findable months later.

USE THIS AS A CHECKLIST

Any AI tool you are about to buy or switch on should pass all five. If it fails three, it is not ready to run your business — it is ready for a small, supervised trial.

Control framing adapted for small businesses. A working checklist, not a published standard.

THE CORE DISTINCTION

Analytics informs a human. AI learns a mapping.

A — ANALYTICS

Turns your data into insight. You decide and act.

B — AI / MACHINE LEARNING

Learns an input → output pattern from examples.

FUNCTION	A — WHAT THE SYSTEM SEES	B — WHAT IT PRODUCES
Sales & marketing	Past purchases, visit frequency, email opens	Which customers are likely to buy again
Bookkeeping	Invoice history, customer, amount, payment terms	Chance this invoice is paid late
Inventory	Sales by item, season, weather, local events	Units to stock next week, by item
Appointments	Booking history, lead time, no-show record	No-show likelihood for a slot
Customer service	The text of an incoming email or review	Its topic and urgency: refund, complaint, lead

If you can write A → B clearly, you understand the use case. Then ask two more questions: what evidence supports the pattern, and what it costs you when B is wrong.

SYSTEM MAP

Classify a system by its role, not by its label.

Seven roles differ by output, who owns the path, and whether it has permission to act.

	ROLE	PRIMARY OUTPUT	WHO CHOOSES THE NEXT STEP?	CAN IT ACT?
1	**ANALYTICS**	Explains evidence	Human	No
2	**PREDICTIVE AI**	Estimates an outcome	Human or rule	No
3	**GENERATIVE AI**	Creates plausible content	Human	No
4	**COPILOT**	Proposes inside a human task	Human	Only through a human
5	**AUTOMATION**	Executes fixed rules	Rule designer	Yes, predefined
6	**AI WORKFLOW**	Uses AI inside a fixed path	Workflow designer	Yes, predefined
7	**AGENT**	Chooses steps and tools	System, within bounds	Yes, bounded

Four questions settle any argument: What is produced? · Who chooses the next step? · Can it act? · Who approves?

HOW TO CLASSIFY · THE 60-SECOND TEST

Three questions sort any tool into one of seven roles.

Ignore the product name and the marketing. Answer in order — then always ask the fourth: who approves?

1 — CAN IT ACT ON ITS OWN?
Send, book, charge or change a record without a person clicking each time?

- **NO** — It hands you something. Go to question 2.
- **YES** — It takes actions in the world. Go to question 3.

2 — IF NO — WHAT DOES IT HAND YOU?
Look at the output, not the brochure.

- **ANALYTICS** — A summary of what already happened
- **PREDICTIVE AI** — An estimate of what will happen
- **GENERATIVE AI** — New content: text, images, replies
- **COPILOT** — Suggestions inside a task you're doing, one step at a time

3 — IF YES — WHO DECIDED THE STEPS?
Before it ran, or while it runs?

- **AUTOMATION** — A person, in advance. Every step is a fixed rule; nothing learns.
- **AI WORKFLOW** — A person, in advance — with an AI step inside the fixed path.
- **AGENT** — The tool itself, while it runs, within limits you set.

Then always ask question 4 — WHO APPROVES? If you cannot name one person, you have found your first job.

CASE 1 OF 7 · CLASSIFY IT

The “AI-powered insights” tab in your Business.

Best sellers, sales by day and hour, average ticket versus last month — the report most owners already glance at.

WHAT THE SYSTEM ACTUALLY DOES

It counts and compares. It adds up completed sales from your point-of-sale system, applies fixed definitions, and shows movement against last week, month or year.

Nothing is predicted. Nothing acts. It hands you evidence and you decide.

WHAT TO WATCH — AND THE CATCH

- Many POS and online-store dashboards now carry an “AI” badge. Check what the tab actually does before paying for an upgrade.
- Reading it is still a skill: it only describes what already happened.
- The catch: a higher average ticket could mean happier customers — or a price rise that quietly cost you your regulars.

WHAT IS PRODUCED?

A summary of sales that already happened

WHO CHOOSES THE NEXT STEP?

You — the owner reading it

CAN IT ACT?

No

WHO APPROVES?

Whoever makes the decision it informs

CLASSIFIES AS

ANALYTICS

? **DISCUSSION**

Many tools sold as “AI-powered” are doing exactly this. What would have to change before this became predictive AI — and would that actually help you, or just give you a longer report?

Illustrative small-business archetype. Check your own POS or e-commerce plan to see what its “AI” feature actually does.

CASE 2 OF 7 · CLASSIFY IT

The forecast that suggests how much stock to order.

A café, shop or salon's software predicts next week's demand by item — and pre-fills the purchase order.

WHAT THE SYSTEM ACTUALLY DOES

It learns from your past sales, the day of the week, the season and local events, then estimates how many of each item you will sell next week.

It estimates. It does not order. The purchase order stays a suggestion until someone approves it.

WHAT TO WATCH — AND THE CATCH

- Forecasts are usually right on average and wrong on the days that matter: a heatwave, a road closure, a viral post.
- It only learns from what you recorded. An item that sold out by noon looks like low demand, not high.
- Note the structure: a prediction becomes a decision the moment you approve every order without reading it.

WHAT IS PRODUCED?

An estimated quantity per item

WHO CHOOSES THE NEXT STEP?

You — or, in practice, the forecast

CAN IT ACT?

Not by design

WHO APPROVES?

You, if you actually read the order

CLASSIFIES AS

PREDICTIVE AI (watch the drift)

? **DISCUSSION**

If you click “approve all” on every suggested order without looking, is this still predictive AI — or has it become automation? What would you check each week to stay the one deciding?

Illustrative small-business archetype. The sold-out point is a general property of learning from sales records, not a claim about any specific product.

CASE 3 OF 7 · CLASSIFY IT

The chatbot that writes your product descriptions.

You ask a general AI assistant to draft listings, a newsletter and replies to online reviews. You edit and post.

WHAT THE SYSTEM ACTUALLY DOES

It generates new text in your chosen tone from a short prompt. Underneath, a large language model predicts the most likely next word, over and over, having learned from vast amounts of text.

The result is fluent by construction. Fluent is not the same as true — and your name goes on it when you post.

WHAT TO WATCH — AND THE CATCH

- Real gains: first drafts in seconds, and help for owners who dread writing.
- The catch: it will happily invent features, materials, certifications or guarantees your product does not have.
- Omissions are quieter than errors. A reply to a complaint can read beautifully and skip the one thing the customer asked.

WHAT IS PRODUCED?

A draft description, email or reply

WHO CHOOSES THE NEXT STEP?

You — you ask and you edit

CAN IT ACT?

No

WHO APPROVES?

You, when you hit publish

CLASSIFIES AS

GENERATIVE AI

? **DISCUSSION**

A listing that says “waterproof” or “organic” when it isn't is a false claim, whoever typed it. What is your two-minute check before AI-written copy goes live — and who does it when you're busy?

Illustrative small-business archetype. The next-word explanation is a simplified summary of how large language models are trained.

CASE 4 OF 7 · CLASSIFY IT

The assistant that suggests expense categories.

Your bookkeeping software reads each bank transaction and proposes a category. You or your bookkeeper accept, change or reject each one.

WHAT THE SYSTEM ACTUALLY DOES

It works inside a task a person is already doing, at that person's pace. You choose when to look, what to accept, and when to close the month.

That is the difference from an AI workflow: there is no hidden path. There is a person doing a job with a helper beside them.

WHAT TO WATCH — AND THE CATCH

- It learns from your past choices, so it gets faster — and repeats your past mistakes just as confidently.
- Wrongly categorised expenses flow straight into your profit figure and your tax return.
- The catch: acceptance rate is not accuracy. A tired owner at 11 p.m. clicks “accept all.”

WHAT IS PRODUCED?

A suggested category for each item

WHO CHOOSES THE NEXT STEP?

You or your bookkeeper, item by item

CAN IT ACT?

Only through you

WHO APPROVES?

You, then your accountant at year-end

CLASSIFIES AS

COPILOT

? **DISCUSSION**

If you start clicking “accept all” every month, has the copilot quietly become an automation? What small sample would you check each month to know the categories are still right?

Illustrative small-business archetype. Many accounting platforms offer suggested categories; check how yours records which suggestions were accepted.

CASE 5 OF 7 · CLASSIFY IT

Automatic payment reminders and late fees.

Your invoicing software emails reminders on a schedule and adds a late fee after 30 days. No AI anywhere in it.

WHAT THE SYSTEM ACTUALLY DOES

An invoice passes its due date and a rule fires: a reminder at 7 days, a firmer one at 14, a late fee at 30.

Every step is a rule someone set. Same input, same output. Nothing is learned and nothing is guessed.

WHAT TO WATCH — AND THE CATCH

- This is the oldest automation in a small business, and it may already be running in your account.
- It acts — it charges your customers money — which is more than most AI tools in this deck are allowed to do.
- Precision trap: a vendor may add an AI "risk score" beside it. Different part, different job. Classify them separately.

WHAT IS PRODUCED?

A reminder, and possibly a fee

WHO CHOOSES THE NEXT STEP?

Whoever set the rules, in advance

CAN IT ACT?

Yes — predefined actions

WHO APPROVES?

Whoever set it up, maybe years ago

CLASSIFIES AS

AUTOMATION

? **DISCUSSION**

This tool can charge your best customer a late fee on a disputed invoice without anyone noticing. When were its rules last reviewed — and do you know who set them up?

Illustrative small-business archetype. Check the reminder and late-fee settings in your own invoicing software.

CASE 6 OF 7 · CLASSIFY IT

The inbox that sorts and routes customer emails.

An AI step reads each incoming email, tags it as a refund, complaint, booking or sales lead, and moves it to the right queue.

WHAT THE SYSTEM ACTUALLY DOES

A designed path: an email arrives, a model reads it and assigns a label, urgent labels alert the owner, the rest wait in a standard folder, and an automatic acknowledgment goes out.

The model sits inside the path. It does not choose the path. A person set up the sequence in advance.

WHAT TO WATCH — AND THE CATCH

- It does not answer the customer. It changes the order in which a person sees them — which matters on a busy day.
- It works well on the common cases it has seen many times before.
- The exception queue is the whole risk: what happens to an angry complaint the model tags as a general question?

WHAT IS PRODUCED?

A label and a routing decision

WHO CHOOSES THE NEXT STEP?

Whoever designed the flow, in advance

CAN IT ACT?

Yes — predefined routing

WHO APPROVES?

You, when you set up the rules

CLASSIFIES AS

AI WORKFLOW

? **DISCUSSION**

A mis-tagged complaint does not produce an error — it produces a quiet email in the wrong folder. Which of your tools could fail silently like this, and how would you ever find out?

Illustrative small-business archetype. Many help-desk and email tools offer AI tagging and routing; the shape of the workflow matters more than the brand.

CASE 7 OF 7 · CLASSIFY IT

The “AI agent” asked to find you a new supplier.

You ask it to compare three packaging suppliers. It plans its own search, picks its sources, and works unattended.

WHAT THE SYSTEM ACTUALLY DOES

No person chose the sequence. The model decided what to search, what to read, what to discard, and when it had enough — then produced a comparison with links.

That one property — choosing its own steps while it runs — makes it an agent. Give it email or a card, and it can act on those choices too.

WHAT TO WATCH — AND THE CATCH

- The vendor of one such research agent states it can sometimes hallucinate facts and often fails to convey uncertainty accurately.
- Its record on a very hard expert benchmark was 26.6% — a record score, not a pass.
- A company describing its own product's weaknesses is the most credible caveat you will ever cite.

WHAT IS PRODUCED?	WHO CHOOSES THE NEXT STEP?	CAN IT ACT?	WHO APPROVES?	CLASSIFIES AS
A researched comparison with links	The tool, within limits you set	Yes — it uses tools, within limits	Nobody, unless you build it in	**AGENT**

? **DISCUSSION**

The answer will look polished and sourced. How would you verify it before you sign a supplier contract — and how long would that take compared with doing the search yourself?

OpenAI, Introducing deep research (2 February 2025), including the vendor's own stated limitations and the Humanity's Last Exam benchmark result.

DEBRIEF · WHAT THE SEVEN CASES TEACH

The architecture decides the answer — not the product label.

1	**ANALYTICS**	POS “AI insights” tab	Describes the past. Read every headline number beside a second one.
2	**PREDICTIVE AI**	Stock-order forecast	A prediction becomes a decision when you approve without reading.
3	**GENERATIVE AI**	AI-written listings	Fluent is not true. Check every claim before your name goes on it.
4	**COPILOT**	Expense categories	Click “accept all” and a copilot silently becomes an automation.
5	**AUTOMATION**	Reminders and late fees	It has the most permission to act and the least attention.
6	**AI WORKFLOW**	Email sorting and routing	It fails silently. Look in the folders nobody checks.
7	**AGENT**	Supplier research agent	Checking its work is the real cost. Budget for it before you buy.

Notice what none of these were solved by: a better AI. Every control that mattered was an owner's decision about evidence, sequence, permission and sign-off.

PRACTICE · QUICK-FIRE · 10 MINUTES

Classify these eight tools. Ignore what they're called.

Use the 60-second test. Write your answer on each card, then compare at your table.

A “AI AGENT” PHONE LINE

Callers press 1 to 5 and hear a recorded answer for each option. Same menu on every call.

YOUR CALL: _____

B SPREADSHEET HELPER

As you type, it suggests the next formula. You press Tab to accept it or keep typing.

YOUR CALL: _____

C “AI-POWERED INSIGHTS”

A dashboard showing last month's best sellers by weekday and hour.

YOUR CALL: _____

D RECEIPT SCANNER

AI reads each receipt photo, picks a category and posts it to your books. Unsure ones go to a review folder.

YOUR CALL: _____

E WRITING ASSISTANT

You ask a chatbot to write this month's newsletter and three new product descriptions.

YOUR CALL: _____

F LATE-PAYER SCORE

Your accounting app scores each customer on how likely their next invoice is to be paid late.

YOUR CALL: _____

G “AI AGENT” FOR QUOTES

You ask it to get three quotes for new signage. It searches, emails local sign shops and compiles the replies.

YOUR CALL: _____

H “SMART” WAITLIST

When a booking is cancelled, the first person on the waitlist automatically gets a text offering the slot.

YOUR CALL: _____

Two of these are sold as “AI agents.” Only one of them is. Label the role, not the brochure.

IN THE WILD · CLASSIFY BEFORE YOU BUY

The AI phone receptionist.

One of the most-pitched tools for small businesses — and a clean illustration of why the setup, not the label, decides the category.

1 CALLER IN

A customer calls after hours and asks, in their own words, for an appointment.

2 MODEL OUTPUT

The AI understands the request and proposes a time from your calendar.

3 YOUR GATE

Decide in the setup: does it book directly, or hold the slot until a person confirms?

WHAT IT IS — AND IS NOT

- If it follows a path you designed — understand, check calendar, book or hand off — it is an AI workflow, not an agent.
- If it can improvise — offer discounts, promise call-backs, answer policy questions — it is closer to an agent and needs tighter limits.
- The value is never missing a call. The risk is a promise nobody checks.

THE QUESTIONS TO ASK THE VENDOR

- What can it say or do without a person — and can I switch each permission off?
- When it isn't sure, where does the call go, and who picks it up?
- Can I read a transcript of every call, and keep it for six months?

A general small-business archetype, not a review of any product. Vendor capability claims are claims, not measured outcomes — test them on your own calls.

THE NUMBER NOBODY FORECASTS

Every AI workflow has an exception queue. That is where the cost lives.

THE CLASSIC ILLUSTRATION

A postal scanner photographs an envelope and a model reads the address. High confidence: a barcode prints and the letter moves. Low confidence: a person types it, then it rejoins the path.

98% read by machine · 841,260,847 images keyed by people in one year

THE VERSION IN YOUR BUSINESS

- Receipts: the ones the scanner cannot read still need you.
- Phone bot: the callers it cannot understand still need a call back.
- Email sorting: the messages it tags wrongly still need finding.

THE BUSINESS-CASE ERROR THIS CAUSES

- “Handles 90% automatically” is sold as hours saved. The remaining 10% is the hardest 10% — and each one takes longer than the average case did before.
- Exception volume is rarely in the vendor's pitch, rarely staffed, and almost never in your own sums.
- Ask for it by name: what share won't it handle, who deals with those, how long each takes, and what happens when they pile upon a busy day?

USPS published address-recognition rate and Remote Encoding Center volume, FY2025. Used here as an illustrative pattern, not a small-business benchmark.

DESIGN HEURISTIC

Autonomy is a design choice — not a magic property.

Start with the least autonomy that delivers value you can measure. A well-checked level 2 beats an unchecked level 4.

1	**ANSWER**	No tools. It produces text; a human does everything else.	*Drafting emails, listings, summaries*
2	**ASSIST**	Read-only tools. It can look things up. A human still acts.	*Looking up orders, reading your FAQ*
3	**EXECUTE**	Predetermined low-risk actions inside explicit limits.	*Booking reminders, waitlist texts*
4	**ORCHESTRATE**	Chooses its own steps — and pauses on anything high-stakes.	*Rarely justified with your money or customers today*

A working heuristic for this workshop — not an industry standard. Being explicit about where a framework comes from is itself the habit being taught.

THE OPERATING PRINCIPLE

AI does not remove the need for analytics. It raises the value of verified evidence.

★ RECONCILE BEFORE YOU AUTOMATE

A confident tool built on the wrong numbers produces a confident wrong answer — faster, at scale, and with your name on it.

Before any automation touches your data: compare the count and one total against your bank statement, POS or books yourself.

GUIDED CASE · SIMULATED

A small business ran a local ad campaign.

METRIC	BEFORE	CAMPAIGN PERIOD	CHANGE
Ad and promotion spend	$4,000	$4,800	+20%
Booked appointments	1,000	1,060	+6%
Net revenue	$50,000	$52,000	+4%
Contribution margin rate	40%	35%	**–5 percentage points**
No-show rate	11%	14%	**+3 percentage points**

? THE DECISION ON THE TABLE

Should the business increase, hold, or reduce ad spend next month — and what would you need to see before deciding either way?

WHAT IS MISSING FROM THIS TABLE

A comparison period or area without ads · service mix · seasonality · new vs. returning customers · an agreed definition of net revenue · whether the extra bookings were coming anyway

All values are simulated for this exercise. No real business, campaign or measured outcome is represented.

AUDIT THE ANSWER · THE HIGHEST-VALUE EXERCISE OF THE DAY

It sounds executive-ready. Is it evidence-ready?

AI-GENERATED DRAFT

“The ad campaign was highly successful: revenue rose 20%, profitability improved, and the business should increase ad spending by 25% next month.”

1

WHICH NUMBER IS WRONG?

Cite the exact cell that proves it. “It seems optimistic” is not an answer.

2

WHICH CAUSAL CLAIM FAILS?

Say precisely what evidence would be required to support it.

3

WHICH ACTION HAS NO RULE?

Name the missing threshold, the missing owner, and who would have to approve.

This is the whole method in one exercise: an output that is fluent, confident, internally coherent and wrong. Fluency is not evidence. Every correction must cite a cell.

THE RULES FLOOR · WHEREVER YOU OPERATE

Four kinds of rules already apply to your AI use.

Orientation only — confirm specific obligations with your attorney, accountant or trade association before acting.

1 Consumer protection & advertising

What your chatbot, listings and ads promise is your promise. Regulators have already acted against businesses making false or inflated claims about AI and with it.

2 Privacy & customer data

Customer names, health details, card data and messages you paste into an AI tool may be covered by privacy law, card-network rules or your own privacy policy.

3 Employment & hiring

Using AI to screen applicants, set schedules or monitor staff can trigger anti-discrimination and notice rules — and some cities and states regulate it directly.

4 Your industry's own rules

Health, finance, food, legal, real estate and childcare all carry sector rules that apply whether a person or a tool did the work.

The common thread: all four ask who is accountable, what the tool relied on, and what you can show afterwards — the same five-control checklist.

THE CONTROL YOU ACTUALLY OWN

Design an approval gate that is real, not decorative.

In a small business, gates fail not because they are missing, but because the owner is too busy to use them.

✓ WHAT A REAL GATE LOOKS LIKE

- A named person — you or a specific employee, not “whoever's free.”
- Shown what they need to disagree: the input, the output, and what the tool did not check.
- Given time to actually look. A gate you clear in two seconds between customers is a rubber stamp.
- Free to say no. If rejecting the tool's output causes hassle, nobody will.
- Kept on record: who approved, what they saw, when.

✕ WHAT A DECORATIVE GATE LOOKS LIKE

- “The tool flags it and we keep an eye on it.” Who, and by when?
- Approvals buried among two hundred identical items a week.
- An approver who cannot see why the tool reached its answer.
- A check on what goes out, but none on the data it ran on.
- Sign-off that is required but never spot-checked.

Test it this way: if this gate had existed in the two chatbot cases, would it have caught anything? If not, it is paperwork, not control.

YOUR NEXT 30 DAYS

Five low-cost actions that change how you run AI.

01

List every tool already running

Include the ones nobody calls AI: reminders, POS insights, email rules, chatbots. Classify each one with the 60-second test.

02

Check one number you rely on

Pick a figure you make decisions on. Trace it to your bank, POS or books. Compare the count and one total. Note what you find.

03

Name the gate on one tool

Pick one tool that touches customers or money. Write down who approves its actions, what they see, and how long they have.

04

Ask for the exception rate

Before you buy the next tool, ask what share it won't handle, who deals with those, and how long each takes.

05

Write one definition down

Take the number your team argues about most — “a sale,” “a new customer.” Agree a definition in writing. Share it.

BRING BACK IN 30 DAYS

One page: your classified tool list, what the number check showed, and one gate you named. The follow-up session works from your material, not ours.

WHERE THIS LEAVES YOU

The question is never “is the AI good?”

It is: what kind of tool is this, what is it running on, what can it do without asking, and whose name is on the approval?

CLASSIFY

Classify by role, not by label.

CHECK

Check the numbers before you automate.

CONTROL

Name the person at the gate.