

# AI Policy Starter Pack for Small Businesses

A practical starter pack for small businesses that want clear AI usage lanes, approval rules, and red-flag guidance before AI tools spread across the team.

Owners, operators, and office leads setting AI usage standards for a growing team

## GUIDE

A practical starter pack for small businesses that want clear AI usage lanes, approval rules, and red-flag guidance before AI tools spread across the team.

## WHAT THIS ASSET COVERS

- Usage lanes for safe delegation, review-required work, and restricted work
- Approval rules for customer-facing, legal, financial, and reputation-sensitive outputs
- Red-flag scenarios that should trigger human review immediately

## USE THIS WHEN

1. The team is already using AI informally with no real operating policy
2. You want a practical governance asset instead of enterprise theater
3. You need a lightweight way to keep customer-facing outputs trustworthy

## WORKING ASSET

Set practical rules for how AI should be used inside a small business before adoption spreads faster than quality control.

## USAGE LANES

Split AI usage into three lanes.

### GREEN LANE

Low-risk work that can move fast:

- draft internal notes
- summarize meetings
- create first-pass outlines
- generate rough ideas for operator review

## **YELLOW LANE**

Useful but review-required work:

- customer-facing email drafts
- FAQ drafts
- review-response suggestions
- estimate follow-up drafts
- content briefs and page outlines

## **RED LANE**

High-risk work that should not be delegated casually:

- legal commitments
- financial approvals
- medical or compliance guidance
- fake proof or invented customer statements
- unsupervised policy or pricing changes

## **APPROVAL RULES**

Use simple approval rules instead of vague “be careful” guidance.

- Any public-facing copy requires a named reviewer.
- Any customer-specific output must be checked for factual accuracy.
- Any content referencing pricing, guarantees, timelines, or sensitive customer data requires human sign-off.
- Any output that could become published proof requires consent and verification.

## **PROMPTING STANDARD**

Require teams to include:

- business context

- audience
- output format
- required facts
- tone constraints
- what must not be changed

That alone improves quality more than most teams expect.

## **HUMAN REVIEW CHECKLIST**

Before approving AI-assisted output, check:

- is it factually correct?
- does it sound like the business?
- does it overstate, invent, or imply unsupported proof?
- does it make commitments the team cannot honor?
- does it protect trust rather than just save time?

## **RED-FLAG SCENARIOS**

Escalate immediately when AI is used for:

- fabricated reviews or testimonials
- policy language the owner has not approved
- legal or compliance claims
- pricing or guarantee changes
- hidden edits to public proof or case examples

## **OWNERSHIP MODEL**

Define:

- who can use AI for what
- who approves customer-facing work
- who updates the policy
- who logs incidents or quality failures

Without ownership, the policy becomes decorative.

## **TRAINING CADENCE**

## WEEKLY

- review one good use case
- review one weak or risky output

## MONTHLY

- update examples
- clarify any lane confusion
- tighten prompts and review guidance

## 30-DAY ROLLOUT

### WEEK 1

- define lanes
- assign reviewers
- document restricted uses

### WEEK 2

- train the team on prompts and review rules
- test real outputs against the checklist

### WEEK 3

- fix weak spots
- clarify edge cases

### WEEK 4

- publish the internal policy
- review adoption quality, not just adoption volume

## DEPLOYMENT NOTES

### HOW STRONG TEAMS ACTUALLY USE THIS ASSET

- Assign one accountable owner instead of letting "AI Policy Starter Pack for Small Businesses" become shared but unmanaged work.
- Use it with owners, operators, and office leads setting ai usage standards for a growing team in a weekly rhythm so the asset drives decisions rather than sitting in a folder.

- Decide in advance what counts as green, watch, and red performance so the team knows when to escalate.
- Capture learnings directly in the document every week so the asset becomes smarter over time instead of resetting to zero.

## HOW TO GET STRONGER OUTPUTS FROM MODERN AI MODELS

- Start with a compact context packet: business type, customer situation, service offered, tone guardrails, and any facts the model must preserve.
- State the deliverable shape up front: channel, word count, required fields, and the exact output format you want back.
- Use variables and clear delimiters so the prompt can be reused safely by staff without rewriting the entire instruction every time.
- Include one strong example when tone and structure matter, then ask for a final answer only rather than hidden reasoning.
- Add a final self-check step for compliance, specificity, and whether the response actually sounds like a real operator wrote it.

## BEST DEPLOYMENT SEQUENCE

- The team is already using AI informally with no real operating policy
- You want a practical governance asset instead of enterprise theater
- You need a lightweight way to keep customer-facing outputs trustworthy

## WHAT SEPARATES A SERIOUS VERSION FROM A BASIC TEMPLATE

- Clear ownership for every step, not generic advice without accountability.
- Targets, thresholds, or decision rules that tell the team what good looks like.
- Specific working components: Usage lanes for safe delegation, review-required work, and restricted work, Approval rules for customer-facing, legal, financial, and reputation-sensitive outputs, Red-flag scenarios that should trigger human review immediately.
- A built-in review cadence so the document becomes part of operations rather than a one-time download.