

Finance AI use policy template

An AI usage policy template a finance team can adopt before any AI agent touches its books: the duties agents may carry, the data they may read, the approvals they need and the review that keeps them honest.

Published 12 October 2026

A finance AI use policy sets the terms on which AI tools and agents may work in a finance team: which duties they may carry, what data they may read, which actions need a person's approval, what they may never do and how the team reviews their work.

The clauses below suit any vendor and any AI tool. They draw on the controls Rever publishes on its [trust](#) and [agents](#) pages, written as obligations your company places on whatever system it uses. Replace each bracketed placeholder with your own value.

How to use the template

- Copy the policy into your controls documentation and name the owner, often the CFO or controller.
- Fill the duties table with the work you plan to hand to agents, and leave out what you do not.
- Set the amount bands with your approval matrix, and have the audit committee or owner approve the policy.
- Review it on the cadence in section 9, and each time you add a vendor or a duty.

The AI usage policy template

1. Purpose and scope

- This policy governs the use of AI tools and AI agents in [Company]'s finance and accounting work, including tools provided by vendors and tools the company builds.
- It applies to every entity, every finance employee and contractor, and every system an AI tool can read or act in.
- The [CFO] owns this policy. Process owners own the SOPs that sit under it.

2. Definitions

- An AI agent is software that carries out a finance duty, such as matching, preparing an entry or drafting a message, with some independence.

- A financial action is any posting, invoice, credit note, payment, payment hold, vendor master change or communication sent outside the company.
- A preparer is the person or agent that prepares or proposes an action. An approver is the authorised person who approves it.
- A material amount is any amount at or above [amount], or any amount the approval matrix marks for dual approval.

3. Permitted duties

DUTY	WHAT AN AI AGENT MAY DO	APPROVAL NEEDED
Bank, subledger and balance reconciliation	Match records, list differences and cite the support	Reviewer signs the reconciliation
Invoice matching before payment	Compare invoice, order, receipt and contract terms; flag exceptions	AP lead decides each exception
Duplicate and price checks	Flag possible duplicates and price differences with history	Approval before any payment hold is placed
Journal entries, accruals and reversals	Prepare entries with calculations and support	Approver within the amount band; dual approval on material amounts
Posting an approved entry	Execute through an authorised workflow and confirm it posted once	Recorded approval before posting
Customer and vendor claims, confirmations and disputes	Prepare the claim or message with evidence	Approval before anything is sent
Vendor bank-detail changes	Flag the change and prepare an independent callback to a contact already on file	A person completes the callback and approves the change
Flux and variance explanations	Draft explanations tied to transactions	Controller reviews before the close pack
Cash forecasting	Build and refresh the forecast from reconciled records	Treasury lead approves any payment or collection action
Answering finance questions	Answer from the books and documents, citing sources and stating uncertainty	None for the answer; the usual approvals for any action that follows

Any duty not in this table needs the policy owner's approval before an agent carries it.

4. Data access

- AI tools connect read-only by default. A tool may write to a company system only through a workflow the policy owner has authorised, after approval.
- Each connection gives access to the records a duty needs and no more.

- Agents do not hold login credentials to company systems. Connections are integrations the company can inspect and revoke.
- No vendor may use company data to train models. Anything a vendor's agents learn about the company stays within the company's tenancy.
- Personal data is handled with the least exposure the duty allows, and the vendor documents how it keeps personal data out of model prompts.
- Original records are kept as received, and calculations and source references stay linked to each finding and workpaper.
- Hosting and data location follow [Company]'s data residency rules, recorded in the vendor contract.

5. Approval thresholds and dual approval

AMOUNT BAND	WHO APPROVES	DUAL APPROVAL
Up to [amount]	[Role]	No
[Amount] to [amount]	[Role]	No
[Amount] and above	[Role] and [Role]	Yes

- Each approval is recorded with the approver's identity, the time and the evidence the approver saw.
- The identity that prepares or proposes an action cannot approve it, whether the preparer is a person or an agent.
- Splitting an action to stay under a threshold is a breach of this policy.

6. What AI agents may never do

- Approve their own work, or any action another agent proposed on their behalf.
- Move money, release a payment or place a payment hold without a recorded approval.
- Change vendor or customer bank details, or use contact details supplied with a change request to confirm one.
- Send any communication outside the company without approval.
- Post to a closed period, or change a posted entry instead of reversing and correcting it.
- Delete, alter or hide audit-trail records, evidence links or approval history.
- Present an assumption as a conclusion, or report a number without its source.
- Count a finding as recovered value before the cash arrives or the credit is applied.

7. SOPs and versioning

- Each duty an agent carries has a written SOP, approved by the process owner before the agent uses it.
- An SOP change takes effect only after approval. The previous version and the history of the runs it governed are kept.
- No agent changes its own procedure in silence. Proposed changes go to the process owner with the reason and the runs that prompted them.
- Each SOP names the lead agent or person, the support, the thresholds and the decisions that stay with people.

8. New duties and earned trust

- A new agent or a new duty starts in observation: the agent works beside the person who owns the duty, and the person acts.
- The process owner widens the agent's role only on a written record of its performance against the SOP, and narrows it the same way.
- The record of each change in an agent's role is available to the policy owner and to the auditor.

9. Review and monitoring

- Monthly: the controller reviews open exceptions, overridden findings and any action taken without the expected approval.
- Quarterly: a person or function independent of the duty tests whether agents followed the approved SOPs and whether the controls worked, and states what it did not test.
- Continuous: someone owns the health of the connections, so stale data and failed syncs appear as issues and never as a clean result.
- Annually, and whenever a vendor or duty is added: the policy owner reviews this policy.

10. Incidents and exceptions

- Anyone who sees an agent act outside this policy reports it to the policy owner the same day.
- The policy owner can pause a duty at once. The duty resumes after the cause is found, the fix is tested and the owner approves.
- Each incident is logged with its cause, its financial effect and the retest result.

11. Approval of this policy

- Prepared by: [Name, role]. Approved by: [Name, role]. Effective date: [date]. Next review: [date].

How Rever maps to the policy

Rever's published controls line up with these clauses. Connections are read-only, direct, deterministic integrations rather than an MCP server or an agent holding credentials, and Rever never writes to your systems: approved journal entries come back as an import file your team loads. Approval thresholds are yours to set. Rever's floors apply in every plan, including Free: nothing leaves Rever without a recorded approval from an authorised approver; the identity that prepares an action cannot approve it; material amounts need dual approval; and each decision is recorded with its inputs, source records and approver.

Sentinel prepares an independent human callback for a vendor bank-detail change, using a contact your team already trusts. Auditor tests whether agents followed approved procedures, Steward watches system health, and each agent follows an SOP your team approved, versioned when it changes. Every new agent starts as an observer; Eve reviews its performance against your approved SOP, and the record of each promotion or step back is yours to read. See rever.ai/trust, [the fifteen agents](#) and [questions before sharing AP data](#).

Questions

What should an AI usage policy template include?

It should set which duties AI tools may carry out in finance, what data they may read, which actions need a person's approval and what they may never do, plus how SOPs change and how often someone independent tests the work. It gives the auditor a control to test and gives your team a rule to point to.

Should an AI agent be allowed to post journal entries?

An agent can prepare entries for a person to approve. If your policy lets an agent post approved entries, a separate person should approve each entry first, material amounts should get dual approval, and the agent should never approve its own entry. Rever's agents never post: an approved entry comes back as an import file your team loads.

How often should the policy be reviewed?

At least once a year, and each time the team adds a vendor or hands a new duty to an agent. Independent testing of the agents' work should run on its own, more frequent cadence.