



Implementing AI in accounting work

Applying AI to controlled accounting work while preserving auditability, evidence, review, and accountable professional judgment.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 29 / EMBER & TEAL

EXECUTIVE WORKING EDITION - 08 OCT 2026

29

AT A GLANCE

Applying AI to controlled accounting work while preserving auditability, evidence, review, and accountable professional judgment.

Accounting work depends on evidence, controls and the judgment of qualified people. AI can support that work by speeding up reading, matching, drafting and exception spotting, but it cannot carry professional responsibility and it must not weaken the audit trail. This playbook is for finance leaders and practitioners who want to use AI to support close, reconciliations, invoice review, variance investigation and reporting, while keeping documented controls, clear review and accountable sign-off. It is operational guidance and does not give accounting, tax, audit or legal advice.

WHO IT IS FOR Controllers, finance managers, accountants and finance operations leads who support close, reconciliations, payables and reporting.

TIME TO USE 60 minutes to read; one finance team session to apply

What you will have at the end

- A work map separating tasks AI can support from judgments reserved for qualified professionals
- Information and access rules suited to financial records
- Request templates for invoice review, reconciliation support, variance analysis and close documentation
- An evidence and audit trail standard for AI-assisted work
- Metrics that show time saved without weakening controls

Contents

01	AI as support for controlled work, not a substitute for it	03
02	Where AI can support and where professional judgment decides	04
03	Information, access and tool rules for financial records	06
04	Invoice review, reconciliation support and variance work	07
05	Keeping auditability when AI is in the workflow	09
06	How reviewers should treat AI-assisted work	11
07	Measuring time saved and control quality together	13
08	Roles, training and routines for the finance team	14
09	Where accounting teams run into trouble, with one illustrative scenario	15
10	A first quarter that produces evidence	17
11	Terms used in this publication	18

EXECUTIVE BRIEF

AI as support for controlled work, not a substitute for it

01



The value of AI in accounting lies in faster preparation and better exception spotting. The control environment, the evidence and the sign-off remain with accountable people.

Finance teams spend much of the close reading documents, matching items, chasing explanations and writing up what they found. These are language and pattern tasks where an assistant can save time. The same work is also governed by controls: segregation of duties, approvals, reconciliations with evidence, and review by someone competent and independent of the preparer.



Introducing AI is therefore a change-management question as much as a technology one. A tool that drafts a variance explanation in seconds is helpful. A tool that invents a plausible explanation, which a busy reviewer accepts, is a control failure. The aim is to gain speed while keeping, or improving, the quality of evidence and review.

- Use AI to prepare, read, match, draft and flag. Keep recording, approving and concluding with people.
- Preserve the audit trail: what was asked, what was used, what was produced, who reviewed and what changed.
- Treat financial records as sensitive: approved tools, restricted access, minimal data.
- Define in advance which outputs need independent review and what evidence that review leaves.
- Involve your internal audit, external auditors or advisers early where AI touches controls over financial reporting.



SCOPE NOTE

This playbook supports qualified professionals and does not provide accounting, tax, audit or legal advice. Accounting standards, regulations and your auditors' expectations vary and change. Check current requirements and your own policies.



Read the sections that follow as a set of options to adapt inside your existing control framework, not a replacement for it. The most useful early step is usually small: pick one preparation task, document the control around it and run it in parallel with the existing method.

WORKING QUESTION

Which control in your close process would be hardest to explain to an auditor if an AI tool were involved?

THE WORK MAP

Where AI can support and where professional judgment decides

02

The line runs between preparation and conclusion. Assistants can help prepare. People conclude, approve and record.

TASK	GOOD FIT FOR AI SUPPORT	HUMAN CHECKPOINT
Reading invoices and contracts for key terms	Strong as a first pass	Preparer verifies against the source; approver sees both
Matching bank or subledger items	Good: proposes matches and flags exceptions	Reviewer examines all exceptions and any match with a difference
Drafting variance explanations	Good as a structure and question list	Preparer confirms drivers from system data and discussion with owners
Close checklists and documentation	Strong: drafts and tidies notes	Preparer owns content; reviewer signs
Flagging unusual entries or duplicates	Good as an extra screen	Accountant investigates each flag; no flag, no assumption of cleanliness
Policy memos and technical positions	Weak: use for outlines only	Qualified professional decides and documents the basis

TASK	GOOD FIT FOR AI SUPPORT	HUMAN CHECKPOINT
Posting, approving, releasing payments, final sign-off	Do not delegate	Authorized person, per your controls



The right-hand column matters more than the middle one. For each task, decide who checks, what they look at and what record they leave. If the answer is vague, the task is not ready for AI support.

Classify by control impact

- Low impact: drafting, formatting, summarizing your own notes. Light review.
- Medium impact: proposing matches, categorizing, preparing analysis that a reviewer will rebuild or sample. Defined review and evidence.
- High impact: anything that changes a recorded amount, an approval or a disclosure. Human performs the action; AI may inform.
- Not permitted: autonomous posting, approval or payment actions without a human decision and your formal control design.



Document this classification and have it agreed by the finance leader, and where relevant by internal audit, so that individual preparers are not left to decide the boundary alone under deadline pressure.

WORKING QUESTION

For your five most time-consuming close tasks, which control impact class applies, and who reviews?



SETTING UP SAFELY

Information, access and tool rules for financial records

03

Financial data includes vendor and bank details, payroll, customer balances and unreleased results. Rules should be explicit and enforced by access design, not memory.

Start from minimum necessary data. Most AI support tasks do not need bank account numbers, personal identifiers or full ledgers. Remove or mask what is not needed. Use only tools that your organization has approved for financial information, with settings you have confirmed on retention and training. Involve your security, privacy and legal contacts, and check requirements that apply to your sector and jurisdictions.

DATA	RULE	PRACTICE
Bank details, payment instructions	Never in unapproved tools	Remove from documents before analysis
Payroll and personal data	Restricted; approved tools only if all	Aggregate or mask
Unreleased results and forecasts	Approved tools, named users	Limit to what the task needs
Vendor and customer names	Approved tools; mask where possible	Use codes in samples
Published statements and public filings	Generally low risk	Cite sources

Access and environment checks

- Tool approved by security and finance leadership for the data class in use
- Named users, with access aligned to existing roles and segregation of duties
- No tool connection that can post, approve or pay without a separate human step
- Outputs saved in the finance document system, not only in a chat history
- Retention of prompts and outputs matches your records policy

- Vendor terms on data use reviewed by the right internal owner

SEGREGATION OF DUTIES

An assistant must not become a back door around separation of roles. If a tool lets one person prepare, approve and release, restrict it. Review integrations with the same care as any new system touching the ledger.



Write a short statement for the close file: which tools are used, for what, by whom and what review applies. It helps new staff, supports audit conversations and keeps the use from drifting.

WORKING QUESTION

Could you show an auditor, today, which AI tools touch financial records and what controls surround them?



TASK PLAYBOOKS

Invoice review, reconciliation support and variance work

04



Each playbook follows the same shape: minimal data in, a structured request, outputs that are checkable, and evidence retained.

Playbook 1: invoice and contract review support

01 Prepare the document

Remove bank details and personal data. Keep the invoice number, dates, amounts, terms and description.



02 Request extraction against a checklist

Ask for specified fields only: supplier, invoice number, date, amount, tax lines as stated, payment terms, purchase order reference.

03 Compare to the source

The preparer checks extracted fields against the original and the purchase order or receipt record.



04 Handle flags

Anything marked unclear goes to a person. Do not let the assistant guess missing values.

REQUEST TEMPLATE: INVOICE EXTRACTION

From the invoice text below, extract only: supplier name, invoice number, invoice date, line descriptions and amounts as printed, total as printed, payment terms as printed, PO reference. Write NOT STATED for anything missing. Do not calculate, correct or infer. List any inconsistency you notice between lines and total, quoting the text.



Playbook 2: reconciliation support

01 Define the rule set

Document match keys, tolerance, timing differences and what counts as a reconciling item before running anything.

02 Use the assistant to propose

Ask for candidate matches and a reason code for each exception, working from masked extracts.



03 Verify in the system or workbook

Recompute totals independently. The reconciliation must foot to the ledger and the statement.

04 Review exceptions

The reviewer examines every unmatched and every match with a difference, and records the resolution.

REQUEST TEMPLATE: RECONCILIATION HELPER

Using the rules below, list proposed matches between the bank extract and the ledger extract, and list every unmatched item with a suggested reason (timing, amount difference, possible duplicate, no counterpart). Do not net or drop items. Show counts and totals for matched, unmatched and each side.



Playbook 3: variance investigation

Provide the variance table and ask the assistant to propose questions and likely data cuts to explore, such as by cost center, vendor, period or entry type. Run those cuts in your system. Draft the explanation only from confirmed drivers and conversations with budget owners. If the assistant offers a cause, treat it as a hypothesis to test, not a finding. 🧐

REQUEST TEMPLATE: VARIANCE

Here is the variance table for operating expenses, current month against budget. Suggest ten questions and data cuts that would explain the largest variances, in order of likely materiality. Do not state causes as facts. For each, name the report or query I should run to test it.

WORKING QUESTION

Which one of these three playbooks has the clearest control around it today, and which needs one written before you start? 🌀

EVIDENCE AND AUDIT TRAIL

Keeping auditability when AI is in the workflow

05

II

If you cannot show what was done, with what, by whom and who reviewed it, the efficiency gain is not worth the audit risk.

Auditability does not require recording every keystroke. It requires enough evidence that a competent independent person can understand how a figure or conclusion was reached and see that someone accountable reviewed it. For AI-assisted steps, that means retaining the inputs used, the request, the output, the verification performed and the reviewer's sign-off, in a place linked to the workpaper. 🌟

EVIDENCE ITEM	WHAT TO KEEP	WHERE
Input	Masked extract or document reference and extraction date	Close folder or workpaper
Request	The template used and any changes	Template library with version
Output	The assistant's response, as received	Workpaper attachment
Verification	Tie-outs, recalculations and source comparisons	Workpaper with preparer initials
Review	Reviewer name, date, items examined, changes made	Workpaper sign-off
Tool and settings	Tool name, approved configuration, date of use	Close file statement

01 Standardize the workpaper

Add an AI assistance section: used or not, tool, purpose, verification, reviewer.

02 Version the templates

Keep a change log so you can show which request produced which output.

03 Sample and test

Reviewers or internal audit sample AI-assisted workpapers each quarter to confirm the evidence exists and verification was real.

04 Brief external parties

Tell your auditors and advisers what you are doing, early and plainly, and take their feedback on evidence expectations.

COMMON MISTAKE

Keeping the assistant's conclusion but not the verification. A reviewer then sees a tidy explanation with no proof. Require the tie-out to be attached, not described.

Reproducibility is a practical test. If the same preparer runs the same documented request on the same masked input, they should get an output of comparable substance, and the verification should confirm it. Outputs will vary in wording, so verification must be anchored to numbers and sources, not to exact text.

WORKING QUESTION

If an auditor sampled last month's AI-assisted workpaper, what proof of verification and review would they find?



VERIFICATION AND REVIEW

How reviewers should treat AI-assisted work

06



Reviewers should expect speed in preparation and be more, not less, rigorous about evidence of checking.

A common failure is complacency. When a draft arrives well-formatted and fast, review gets lighter. Reviewers should do the opposite: ask what the preparer verified, and sample against source. The failure patterns are predictable: plausible but invented explanations, dropped items in long extracts, arithmetic that is slightly off, misread dates and currency, and a confident tone over thin evidence.

Reviewer checklist for AI-assisted work

- Preparer states whether and how AI was used
- Totals foot to the ledger, statement or source, with the tie-out attached
- All exceptions and flagged items have a documented resolution
- Sampled items agree to original documents
- Explanations are supported by system data or owner confirmation, not only by narrative
- Anything judgmental is traced to a named qualified person's conclusion
- Dates, currency and period cut-off are correct



FAILURE PATTERN	EXAMPLE	DETECTION
Dropped items	Extract of 400 lines summarized from 380	Row counts and control totals before and after
Invented rationale	Variance attributed to a vendor increase that did not occur	Trace to the transaction detail



FAILURE PATTERN	EXAMPLE	DETECTION
Misread figures	Decimal or thousand separators confused	Recompute from the source
Overconfident classification	Item treated as routine that needs judgment	Reviewer applies policy criteria
False reassurance	No duplicates found because the check was shallow	Run an independent duplicate test

NEVER DELEGATE

Determining accounting treatment, assessing materiality, concluding on estimates, approving journals, releasing payments, signing off reconciliations and any communication of conclusions to auditors, regulators or boards. These belong to accountable, qualified people.

Rebuild one thing independently each cycle. Pick a figure, recompute it by a different route and compare. Over time, this tells you how far you can trust each type of AI-assisted step, and it keeps your reviewers' own skills sharp.

WORKING QUESTION

What does your review do today that would catch an invented explanation that reads well?

MEASURING RESULTS

Measuring time saved and control quality together

07

A faster close that carries more corrections or weaker evidence is a step backward. Track both sides.

Baseline the tasks before changing them. For two cycles, record the hours spent on selected tasks, such as invoice coding checks, bank reconciliation preparation or variance commentary, and the number of review points raised. Then apply the AI-supported method on one task at a time and record the same figures, including verification time. Where possible, run the new method in parallel with the old for a cycle and compare outputs.

METRIC	HOW TO MEASURE	HEALTHY SIGNAL
Preparation hours per task	Time log including verification	Falling
Review points per workpaper	Count of reviewer comments requiring change	Flat or falling
Post-close adjustments	Entries or corrections after close due to preparation error	Not rising
Evidence completeness	Share of sampled workpapers with full AI section and tie-out	Near 100 percent
Exception turnaround	Days from flagged to resolved	Falling
Policy breaches	Use outside approved tools or rules	Zero, reported openly

ILLUSTRATIVE NUMBERS ONLY

If bank reconciliation preparation took 10 hours a month and now takes 7 including verification, the saving is 3 hours a month for that task. These numbers are illustrative. Use your own data, and count the extra evidence work.

Share the results with your auditors and internal audit as well as finance leadership. Transparency about both benefits and incidents builds the trust that lets you extend use. Hold

off on extending to a new task until the quality metrics for the current one have been stable for at least two cycles.


WORKING QUESTION

Which quality metric would tell you first that speed was being bought with weaker control?


TEAM ROUTINES

Roles, training and routines for the finance team

08

Clear ownership and short, regular routines keep AI use consistent across preparers and periods.



- Process owner: one named person for each AI-supported task, responsible for the template, the control and the evidence standard.
- Template library: versioned request templates with the verification step attached to each.
- Pre-close briefing: five minutes on which tasks use AI support this cycle and any changes.
- Post-close retrospective: what saved time, what was caught in review, what to change.
- Quarterly control review: sample workpapers, test evidence and update the control description.

Training

Train on practical cases using masked data from past periods. Include a planted error in each exercise, such as a dropped line or an invented explanation, so participants practice catching it. Teach reviewers as well as preparers. Emphasize that professional skepticism applies to AI output exactly as it does to any representation from a third party.



01 Pilot

Choose one task, one preparer and one reviewer for two cycles, running in parallel with the existing method.

02 Document

Write the control description, template, evidence requirements and reviewer expectations.

03 Assess

Review metrics and incidents with finance leadership and, where relevant, internal audit.

04 Extend

Add tasks one at a time, repeating the same steps.

CULTURE POINT

Make it safe to say that AI was wrong. If preparers fear blame, they will hide errors and the reviewer will never learn where the weak points are. Treat caught errors as the system working.

WORKING QUESTION

Who in your team is the named owner for the first AI-supported task, and what is the control description?

PITFALLS

Where accounting teams run into trouble, with one illustrative scenario

09

Problems arise mainly from weakened evidence, blurred roles and unexamined confidence in output.

- Allowing an integration to post or approve without a separate human step.
- Accepting variance narratives that were never traced to transactions.
- Pasting vendor bank details or personal data into unapproved tools.

- Leaving the assistant's output unsaved, so the workpaper cannot be reproduced.
- Giving the assistant authority over materiality or accounting treatment.
- Not telling auditors until a question arrives.



ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

A fictional company's accountant uses an assistant to explain a large rise in freight expense. The explanation cites a rate increase. The reviewer, following the checklist, asks for the transaction cut and finds the real cause: a duplicated vendor invoice. The narrative was plausible and wrong. Because the workpaper required a tie-out, the duplicate was caught before close.



Notice what saved the day: a process requirement, not a smarter prompt. Hypotheses from an assistant are useful for deciding where to look. They become findings only when the data confirms them.

If something goes wrong

- 01 Stop and preserve**
Pause the use, keep the inputs, outputs and logs.
- 02 Assess impact**
Determine which periods, workpapers and figures may be affected, with the controller.
- 03 Correct and disclose internally**
Fix errors through your normal correction process and inform the right stakeholders, including auditors where appropriate.
- 04 Strengthen the control**
Update the template, evidence requirement and training based on what happened.



Follow your organization's policies and consult qualified advisers on any matter that touches reported results or regulatory obligations.



WORKING QUESTION

What is the one failure above that your current process would be slowest to detect?

30 / 60 / 90

A first quarter that produces evidence

10



Pilot one preparation task in parallel with your existing method, document the control, and extend only when quality metrics hold.

First 30 days: ground it

- Classify close tasks by control impact and agree the boundary with the finance leader
- Confirm approved tools, access and data rules with security and legal contacts
- Baseline hours and review points for two candidate tasks
- Brief internal audit on the plan



Days 31 to 60: test it

- Pilot one task in parallel with the existing method for a full cycle
- Add the AI assistance section and tie-out requirement to the workpaper
- Train preparers and reviewers with planted-error exercises
- Compare time and review points to the baseline



Days 61 to 90: decide

- Sample workpapers to test evidence completeness
- Update the control description and share results with auditors
- Extend to a second task only if quality metrics are stable
- Set a quarterly control review



GLOSSARY

Terms used in this publication

11

TERM	MEANING
Audit trail	The recorded path from source documents to a reported figure, showing who did what and when.
Segregation of duties	Dividing preparation, approval and release among different people to reduce error and fraud risk.
Workpaper	Documentation of work performed, evidence examined and conclusions reached.
Tie-out	Agreeing a figure to its source or to another independent figure and explaining any difference.
Reconciling item	A known difference between two records, such as a timing difference, that explains why they do not agree.
Variance analysis	Comparing actual results to budget, forecast or prior period and investigating the differences.
Materiality	The size or nature of an item that could influence decisions of users of the information, judged by qualified people.
Control description	A written statement of what a control does, who performs it and what evidence it leaves.
Professional skepticism	A questioning attitude toward evidence and representations, applied equally to AI output.
Cut-off	The boundary that decides which period a transaction belongs to.
Control total	A figure from the source used to prove nothing was dropped or duplicated.
Parallel run	Running a new method alongside the existing one for a period to compare results.

A note on use

This playbook supports qualified professionals and is not accounting, tax, audit, financial or legal advice. Check current standards, regulations and your own policies.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

