



Implementing AI in financial analysis

Applying AI to financial-analysis workflows with explicit data controls, transparent assumptions, and expert review.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 36 / EMBER & SKY

EXECUTIVE WORKING EDITION - 08 OCT 2026

36

AT A GLANCE

Applying AI to financial-analysis workflows with explicit data controls, transparent assumptions, and expert review.

Finance teams spend large parts of each cycle explaining numbers rather than producing them: variance commentary, forecast narratives, model checks and management packs. AI can speed the drafting, structuring and cross-checking around that work, if the data is controlled and every figure is verified. This playbook covers variance commentary, forecasting support, model review, scenario framing, data validation and audit trail. It supports qualified analysts and is not investment or financial advice, and AI never replaces their review.

WHO IT IS FOR FP&A analysts, finance business partners, controllers and finance leaders who prepare, review and rely on financial analysis.

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

What you will have at the end

- A work map of finance analysis tasks with good fits for AI and the reviewer checkpoint for each
- Data controls for ledgers, forecasts, customer and employee-related figures
- Request templates for variance commentary, forecasting support, model review, scenarios and validation
- A verification routine, an audit trail and a baseline for measuring time saved and quality

Contents

01	AI for the explaining, not the deciding	03
02	Where AI fits in financial analysis	04
03	Data controls, approved tools and the audit trail	05
04	Variance commentary and forecasting support	07
05	Model review, scenario framing and data validation	09
06	Verification and review habits	11
07	Measuring time saved and quality	12
08	Team routines and sharing what works	13
09	Pitfalls and a worked scenario	15
10	A first quarter that produces evidence	16
11	Terms used in this publication	17

EXECUTIVE BRIEF

AI for the explaining, not the deciding

01



Use AI to draft, structure and challenge. The numbers, the judgments and the sign-off stay with qualified people.

Month-end and planning cycles share a pattern. The numbers arrive, then analysts spend days turning them into explanations: why revenue missed, why a cost line moved, what the forecast now says and what could change it. This is skilled work, but a large part is assembling facts and writing them in a consistent form. That part is a sensible target for AI help.

Finance is also a field where errors carry weight. A language model can misread a table, apply the wrong sign, invent a reason for a variance or produce arithmetic that does not foot. It has no knowledge of your business events unless you provide them, and it cannot take responsibility for a figure. These limits define the playbook: AI helps with words and structure, and every figure is traced to its source by a person.

- Use AI to draft variance commentary from numbers and drivers you supply.
- Use AI to structure forecast narratives and to list questions a challenger would ask.
- Use AI to review a model's logic for inconsistencies, never to replace the model.
- Use AI to frame scenarios and the assumptions each rests on.
- Do not use AI to produce investment recommendations, tax or accounting positions, or to give financial advice.

**SCOPE NOTE**

This playbook supports qualified analysts in internal analysis. It does not provide investment, tax, accounting or financial advice. Check applicable reporting, audit and regulatory requirements with qualified professionals.



WORKING QUESTION

Which part of your last close pack was mostly assembling known facts into sentences?

THE WORK MAP

Where AI fits in financial analysis

02

Separate drafting and checking tasks from judgment and attestation tasks.

List your recurring analysis tasks across the cycle, from data pull to management pack. For each, ask how much is writing and structuring, and what the consequence is if an error slips through. Internal commentary that a controller reviews has a different risk from a figure that reaches lenders, auditors or the public.

TASK	GOOD FIT FOR AI	HUMAN CHECKPOINT
Variance commentary	Draft explanations from supplied variances and drivers	Analyst confirms each driver with the business owner
Forecasting support	Structure assumptions, list risks, draft narrative	Forecast owner sets and signs assumptions and numbers
Model review	Spot inconsistent formulas, hard-codes, unclear logic in described structure	Model owner and independent reviewer test the model
Scenario framing	Propose scenarios, triggers and sensitivities to examine	Finance lead chooses scenarios and ranges
Data validation	Suggest reasonableness checks, reconciliations and test cases	Analyst runs checks and resolves differences
Management pack drafting	Draft summaries, headlines and talking points	Controller checks all figures before release

TASK	GOOD FIT FOR AI	HUMAN CHECKPOINT
Audit trail and documentation	Draft change logs and methodology notes	Preparer verifies against actual work done

RED ZONE

Accounting treatments, valuation conclusions, tax positions, covenant calculations and anything that will be attested or filed are decided and signed by qualified people. AI may prepare questions, never conclusions.

WORKING QUESTION

Which outputs of your team reach people outside the company, and are they clearly in the red zone?

SETTING UP SAFELY

Data controls, approved tools and the audit trail

03

Financial data is sensitive and often regulated. Classify it, restrict where it goes and record how AI was used.

Finance data includes unreleased results, which may be material non-public information, plus payroll, customer and supplier details. Handling mistakes can have legal and reputational consequences. Write a one-page data rule that classifies what you hold and states which tools may touch each class. Confirm requirements with your compliance, legal and audit teams, and check current rules that apply in your sector.

DATA CLASS	EXAMPLES	RULE
Public	Published results, market data from public sources	Any approved tool

DATA CLASS	EXAMPLES	RULE
Internal aggregated	Departmental budgets, monthly variances without names	Approved company tools only
Restricted	Unreleased results, forecasts, board packs, covenant data	Only approved secure environments with explicit sign-off
Personal or counterpart	Payroll, individual compensation, customer or supplier detail	Do not paste; use aggregates or labels

Controls to put in place

- Approved tool list with a named owner and the date terms were checked.
- Confirmation of whether inputs are retained or used to train models under your plan.
- Rule that unreleased results and forecasts are not used in general tools.
- Aggregation or labelling before any analysis support is requested.
- A log of AI-assisted outputs: prompt, date, tool, preparer and reviewer.

The audit trail

Auditors and reviewers will ask how a figure or explanation was produced. Keep, for each AI-assisted deliverable, the data version used, the request, the output as received, the edits made and the name of the person who verified it. This costs a few minutes and answers most later questions.

WORKING QUESTION

If your external auditor asked how AI was used in this quarter's pack, could you show them in an hour?

TASK PLAYBOOKS: VARIANCE AND FORECASTING

Variance commentary and forecasting support

04



These are the highest-volume tasks. Give the tool facts, ask for structure, then confirm every reason with the business.

Playbook 1: Variance commentary

- 01 **Prepare the variance table**
Compute variances in your own system. Provide actual, budget or prior, difference and percentage, with units and period stated.
- 02 **Add the known drivers**
List what you have confirmed: volume, price, mix, timing, one-offs, with the source of each.
- 03 **Draft**
Ask for commentary using only the supplied numbers and drivers, marking any line where the driver is unknown.
- 04 **Verify and confirm**
Recompute each figure quoted. Check the direction of each variance, favorable or unfavorable. Ask the budget holder to confirm each reason.



REQUEST TEMPLATE: VARIANCE COMMENTARY

Write commentary for a management pack. Use only the data below. For each line over the threshold, state the variance in amount and percent, say whether it is favorable or unfavorable, and give the driver from the list. If no driver is given, write Driver to be confirmed. Do not speculate on causes. Table and drivers: [paste, labels not names]



Playbook 2: Forecasting support

Treat the forecast model as yours and use AI around it. Good uses include turning a set of assumptions into a clear narrative, listing what has changed since the last forecast, and

generating the questions a skeptical reviewer would ask. Poor uses include asking the tool to predict revenue or to supply growth rates from general knowledge. Its figures will lack a source, and you cannot defend them.



01 Document assumptions

List each driver, its value, source and owner.

02 Ask for a challenge

Request the questions a reviewer would ask and the assumptions most likely to be wrong, with reasons.

03 Draft the narrative

Have the tool write the forecast story from the assumptions and the model outputs you supply.



04 Reconcile

Confirm every number in the narrative against the model and every change against the prior version.

REQUEST TEMPLATE: FORECAST CHALLENGE

Below are our forecast assumptions and the change from last quarter. Act as a skeptical finance reviewer. List the ten questions you would ask, the three assumptions most sensitive to error, and any assumption that conflicts with another. Do not provide forecast numbers. Assumptions: [paste]



WORKING QUESTION

In your last forecast, which assumption had no named owner or source?



TASK PLAYBOOKS: MODELS, SCENARIOS AND VALIDATION

Model review, scenario framing and data validation

05



Use AI as a second reader for logic and a source of ideas for stress testing, never as the calculation engine of record.

Playbook 3: Model review

Describe the structure of the model, or paste anonymized formulas, and ask for a review against a list of common issues: hard-coded numbers inside formulas, inconsistent formulas across a row, sign errors, circular references, mixed units, missing links between inputs and outputs, and unclear labelling. The tool will find some and miss others, so its output feeds your review rather than replacing it. Run your own integrity checks and, for material models, an independent review.

REQUEST TEMPLATE: MODEL REVIEW

Review the following model description and formulas as a model auditor. List: hard-coded values inside formulas; formulas that differ from neighbors in the same row or column; possible sign or unit errors; inputs that do not flow to outputs; labels that are ambiguous. For each, give the cell or section and why it matters. Do not rewrite the model. Material: [paste]



Playbook 4: Scenario framing

01 Define the question

State what decision the scenarios inform and the time horizon.

02 Generate candidates

Ask for scenarios, the triggers that would indicate each and the variables most affected.

03 Select and range

The finance lead selects scenarios and sets the ranges, based on business knowledge.



04 Run in the model

Calculate in the model of record, and document assumptions for each scenario.



REQUEST TEMPLATE: SCENARIOS

We sell subscriptions to mid-sized firms and are planning next year. Suggest six scenarios that could materially change revenue or cost, such as customer churn rising, a key supplier price change or a delayed product launch. For each give the trigger to watch, the lines affected and one mitigation to consider. Do not give numbers.

Playbook 5: Data validation



01 List expected relationships

Ask for reasonableness checks: totals that must tie, ratios with normal ranges, sign rules, period-over-period tolerances.

02 Run the checks yourself

Implement them in your own tools against the actual data.

03 Investigate breaks

Resolve each exception to a cause: timing, mapping, error or genuine change.

04 Record results

File the checks and outcomes with the pack for the audit trail.



WORKING QUESTION

What check, if it had existed, would have caught your last reporting error before the pack went out?



VERIFICATION

Verification and review habits

06



Finance checks must be mechanical and repeatable. Do not rely on the tool's own claims of accuracy.

Language models are weaker at arithmetic and at consistent handling of signs and units than at prose. They can state a total that does not foot or reverse a favorable and unfavorable label with complete fluency. Therefore every number in an AI-assisted draft is traced to its source and recomputed by a person or by your own tools, not by the same model.



Before an AI-assisted analysis is shared

- Every figure matches the source report or model, to the stated precision.
- Totals foot and percentages recompute.
- Favorable or unfavorable direction is right for each line.
- Each stated reason is supported by a driver you provided or confirmed with the owner.
- Period, entity and currency labels are correct.
- No restricted or personal data was entered into an unapproved tool.
- A reviewer other than the preparer has signed off, and the audit-trail record is complete.



Common failure modes

- Invented explanations: a plausible reason for a variance nobody confirmed.
- Wrong sign or direction in favorable and unfavorable language.
- Rounding differences presented as precise statements.
- Mixing periods, entities or currencies when data from several tables is supplied.
- Narrative that sounds certain about a forecast that is only an assumption.
- Citing external statistics without a source.



FOOTING TEST

Ask a colleague to take the commentary only and rebuild the table from it. If numbers do not rebuild, something is wrong in the text.

WORKING QUESTION

Who reviews the numbers in your pack independently of the person who drafted the text?

MEASURING

Measuring time saved and quality

07

Count hours across the whole cycle and track errors found at review and after release.

Choose two or three recurring outputs, such as monthly variance commentary and the quarterly forecast narrative. Record how long the last few took, from data ready to pack issued. Then run the AI-assisted process and record the same, including verification time. Also record review findings, because speed that comes with more restated numbers is a loss.

METRIC	HOW TO CAPTURE	SIGNAL
Hours from data ready to pack issued	Preparer time log	Net saving including checks
Review comments per pack	Reviewer count by type	Draft quality
Post-issue corrections	Restatements or amendments after release	Should not rise
Share of lines with confirmed driver	Count against commentary lines	Evidence discipline

METRIC	HOW TO CAPTURE	SIGNAL
Validation breaks found	Count exceptions before issue	Value of the checks
Reviewer confidence	Short monthly rating	Over- or under-trust



ILLUSTRATIVE EXAMPLE

If commentary takes 8 hours today and 4 hours with AI help plus 1.5 hours of extra verification, the saving is 2.5 hours. If corrections after issue also rise, the gain is false. These figures are illustrative.

Use two or three cycles before drawing conclusions. One good month can be luck.

WORKING QUESTION

What was your correction count last quarter, and would you know if it rose?



TEAM ROUTINES

Team routines and sharing what works

08



Standard templates, a clear reviewer role and a short monthly review keep practice consistent.

Finance teams value consistency, and AI use should be no different. Keep one shared set of request templates, one checklist and one log. Make the review step part of the close calendar so it is not squeezed when deadlines tighten.



01 Standardize the inputs

Agree the variance table format, driver list format and labelling so requests are the same each time.

02 Fix the roles

Name the preparer, the independent reviewer and the owner of the tool list and templates.

03 Put review on the calendar

Add a verification step to the close timetable with time allowed.

04 Meet monthly

Thirty minutes to share what worked, what slipped through and what to change.

05 Coach judgment

Let juniors draft their own commentary first sometimes, so analytic skill continues to grow.

- Share templates, not data. Examples shown to the team use labels and invented figures.
- Keep a log of near misses, such as a reversed variance caught at review.
- Brief the audit and compliance teams on how you use AI and what controls apply.
- Review the rules when tools, contracts or regulations change.

Playbook 6: Management pack drafting

Once variances, forecast outputs and validation checks are done, the pack itself is a drafting task: a headline page, key messages, risks and opportunities, and talking points for the presenter. Give the tool only verified figures and confirmed drivers, tell it the reader and the length, and ask for a structure that leads with what changed and what decisions are needed. Then recheck every figure in the final text against the pack tables, because edits made late are where numbers drift.

REQUEST TEMPLATE: PACK SUMMARY

Write a half-page executive summary for the monthly management pack. Use only the figures and confirmed drivers below. Lead with the three most important changes versus budget, then risks and opportunities, then decisions needed. Use plain language, state units and period, and write Driver to be confirmed where none is given. Do not forecast or advise. Data: [paste, labels not names]

- Ask the tool to list questions an executive is likely to raise, and prepare answers in advance.
- Keep a fixed glossary of metric definitions so terms mean the same thing every month.
- Have the presenter read the summary aloud once, to catch claims they cannot support.
- Store the final pack with the log entry so the version issued is the version recorded.

WORKING QUESTION

Is verification a named step on your close calendar, or something people fit in if time allows?



PITFALLS

Pitfalls and a worked scenario

09



The common failures are invented explanations, unverified numbers and loose handling of restricted data.

- Asking the tool to explain a variance without giving it the drivers.
- Pasting unreleased results into a general-purpose tool.
- Accepting arithmetic from a language model.
- Presenting scenario text as a forecast.
- Skipping independent review under time pressure.
- Letting AI wording imply certainty or advice that the analyst would not give.



ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

An analyst asks an assistant to explain a rise in selling costs. With no drivers supplied, it suggests higher freight rates and writes it confidently. The real cause was a timing shift in a marketing invoice. The budget holder, asked to confirm the reason, corrects it before the pack goes to the executive team. The control that worked was a rule: no driver, no explanation.

The scenario is invented, yet the pattern is typical: when asked to explain without evidence, the tool produces something plausible. Supplying drivers, and confirming them with owners, turns that weakness into a managed step.



Halden & Company offers free assessments at halden.company/assessments to help finance teams find safe uses of AI, and you can reach the team at halden.company/contact. This material is not investment or financial advice.

**WORKING QUESTION**

Which explanation in your next pack is currently supported only by plausibility?

30 / 60 / 90

A first quarter that produces evidence

10



Start with variance commentary on aggregated data, build the audit trail from day one and extend only as checks prove reliable.

First 30 days: ground it

- Write the data classification and approved tool rules, with compliance and audit input.
- Map finance tasks into green, amber and red zones.
- Baseline time and review findings for one recurring output.
- Create the variance commentary template and the verification checklist.

Days 31 to 60: test it

- Pilot variance commentary on aggregated data with independent review.
- Start the AI-assisted output log for the audit trail.
- Add forecast challenge and model review templates.
- Put the verification step on the close calendar.



Days 61 to 90: decide

- Compare pilot with baseline on hours, review comments and post-issue corrections.
- Add scenario framing and validation checks where results justify.
- Brief audit and compliance on practice and controls.
- Decide what to scale, adjust or stop and record why.



GLOSSARY

Terms used in this publication

11



TERM	MEANING
Variance analysis	Comparing actual results to budget, forecast or prior periods and explaining the differences.
FP&A	Financial planning and analysis: the function that budgets, forecasts and analyzes performance.
Driver	A factor, such as volume, price or mix, that explains a change in a financial figure.
Favorable and unfavorable	Labels showing whether a variance improves or worsens profit relative to the comparison.
Scenario	A coherent set of assumptions describing one possible future for analysis.
Sensitivity	How much an output changes when a single input changes.
Foot	To confirm that the components of a table add up to the stated total.
Model of record	The controlled model whose numbers are used for official analysis.
Audit trail	A record of data, requests, outputs, edits and reviewers behind a deliverable.
Reasonableness check	A test that a figure falls within an expected range or relationship.
Material non-public information	Unreleased information that could influence an investor's decision and must be tightly controlled.
Hallucination	A fluent but false or unsupported output from an AI system.



A note on use

General guidance for qualified finance professionals. It is not investment, tax, accounting or financial advice. Check applicable reporting, audit and regulatory requirements.



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.

