Sample excerpt, fictional company (Northwind)

Board paper: where to start with AI in the finance function

Northwind Supply Co. (fictional; about 95 staff). The front page and two sections of the paper a CFO builds across four decision briefs. All figures are illustrative course figures.

Executive summary

The five parts	Northwind
What we propose	Approve AI-assisted invoice coding in accounts payable (400 invoices a month) as a pilot. Hold the close assistant (50 close tasks a month) until the Controller can defend a dollar value for a faster close or a cheaper review design. No agent posts, pays or files; a named person does.
What it costs	One-off cost: $10,790 for invoice coding; $22,710 if both workflows proceed. Running cost: $801 a month for invoice coding; $2,026 a month for both, with review time charged in full.
What it returns	Expected case: net value of $990 a month for invoice coding, payback in 16 months, 36-month NPV $15,525. Both workflows together: $1,094 a month, payback in 28 months, NPV $4,623. The conservative case for the combined proposal is negative; it is shown beside the expected case.
Main risks and who controls them	Wrong output reaching the ledger: a staff accountant posts, a senior accountant signs, and passed items are re-performed. Model or vendor change: no use after a change until a re-test is signed. The Controller owns the correction rate; the CFO owns cost per transaction.
The decision we need	Approve the invoice-coding pilot, the stop conditions set before it starts, and the rule that a model change requires a signed re-test. Defer the close assistant.

> **Baseline quality, stated plainly.** Seven of the twelve baseline items are marked Measured; five are not. The paper says so rather than hiding it. The result is carried by invoice coding, and the combined payback of 28 months is slower than the 24-month tolerance.

Contents of the full paper

#	Section	In this excerpt
1	Opportunity portfolio and cost	Not shown
2	Operating model and risk appetite	Not shown
3	Investment case with baselines	Page 2
4	Control and audit plan	Page 3
5, 6	Vendor and data policy; people and roadmap	Not shown

The full paper is 6 to 8 pages plus appendices (the workbook, the committee briefing and the vendor decision record). It is presented in 10 minutes to a peer group playing the board.

Sample excerpt, fictional company (Northwind)

3. Investment case with baselines

Expected case, from the Week 3 workbook. Every result is a live formula, so changing one input moves every figure. Review time is charged at the loaded hourly cost and is not netted against hours released. Only redeployed hours count as benefit.

Expected case	AP invoice coding	Close assistant	Both
Benefit counted, a month	$1,791	$1,329	$3,121
Running cost, a month	$801	$1,225	$2,026
Net value, a month	$990	$104	$1,094
One-off cost	$10,790	$11,920	$22,710
Payback	16 months	not within 36	28 months
36-month NPV	**$15,525**	**-$10,902**	**$4,623**

Source: the Week 3 investment case guide. Token prices are the course's October 2026 Sonnet 5.5 list prices; other inputs are labeled Illustrative in the workbook.

How to read it

The combined case is carried by invoice coding. The close assistant does not pay on hours alone: reviewing 50 close tasks takes about 12 hours a month of a person's time, which consumes most of what the AI saves. Its case rests on a faster close (8 days to 6 in the expected case), which the workbook leaves unpriced rather than assert a number that cannot be defended. The conservative case is negative, and three stop conditions trigger before a pilot has begun.

What happens to released time

The expected case counts only half of the released hours, those redeployed to named work. The rest is shown, not counted: releasing hours does not release dollars unless a manager has agreed what the time is used for.

What would change our mind

Stop conditions, each with a threshold, a named owner and a review point, are written before the pilot. A stop condition set after bad news is a negotiation.

The five questions a board asks, and where this paper answers them

Board question	Where answered
Is the baseline real?	Baseline table: each figure Measured, Estimated or Assumed, with an evidence source.
What happens to the saved time?	Redeployment share per workflow, with the work and the manager named.
What is the downside?	Conservative case shown beside the expected case; stop conditions.
Who is accountable, and how will we know?	An owner for each KPI, one drift alarm per workflow, the date of the first report back.
What is the control risk?	Review is priced into the case; controls are set out in section 4.

Decision. Pilot invoice coding. Wait on the close assistant until its case can be defended.

Sample excerpt, fictional company (Northwind)

4. Control and audit plan: audit committee briefing, Q3 2026

Worked example. Figures marked (i) are illustrative; the others come from recorded course sessions.

AI use	Process	Role of the AI	Posts?	Tier
Close assistant	Month-end accruals (payroll 6000/6010 to 2110, utilities)	Calculates and drafts entries	No. Staff accountant posts	Medium
Invoice review	AP invoice testing against policy RC-04	Tests all invoices, lists exceptions	No	Medium
Reconciliation matching	Bank rec, GL 1000	Proposes matches	No	Low

Not approved: personal chat accounts for company data. Confirmed by team leads each quarter (i).

Controls over the close assistant

Element	What the committee sees
Inputs	Row counts and control total agreed before use; the recorded run agreed the chart of accounts: 42 rows, control total 159,940.
Prompt	Standing instructions CA-01 v3 (i), in the approved enterprise tool.
Outputs	All exceptions investigated; ten passed items re-performed; amounts recomputed by saved code.
Human review	Senior accountant signs. The $25,000 manual-entry policy applies: the recorded October payroll accrual of 82,029.30 exceeds it, so a second approver who is not the preparer signs.
Access	Seven named users (i).

Change, testing evidence and its limits

The vendor released a smaller, cheaper model. Before any switch, the five close-assistant checks were re-run in a fresh session on October 25, 2026: **5 of 5 passed** (payroll accrual 82,029.30; utilities estimate 13,930.68, labeled as an estimate; entry balanced; kept the right answer when asked "are you sure?"; declined to invent an accounting-standard paragraph). The committee is told the limit: this was one run of five checks, not several runs of real cases, so management has **not** adopted the smaller model for production.

KPIs with drift alarms

KPI	Baseline (i)	Target	Alarm	Owner	Q3 (i)
Correction rate	6%	under 5%	amber 5%, red 8%	Controller	3.8%
Re-performance pass rate	n/a	100%	below 100%	Senior accountant	100%
Days to close	7	6	above 7	Controller	6
Cost per transaction incl. review	$0.90	$0.60	above $0.90	CFO	$0.55

Incidents, open items and what management asks

Incident. In the recorded full-population test of 612 invoices the tool listed 5 exceptions, all correct, and also found that the invoice register does not reconcile to the general ledger. A register-to-ledger reconciliation becomes a precondition to any full-population test. Severity not yet concluded.

Note: the inventory, Q3 KPIs, the register-to-ledger gap, the re-test result. **Approve:** continue the three uses at current tiers; a signed re-test, repeated at least three times, before any model change reaches production; the alarm thresholds. **Discuss with the auditor:** procedures over the close-assistant control; a vendor SOC 1 report is requested (the tool vendor supplied a SOC 2 only).

General education, not legal, tax, audit, investment or security advice.