

eFuturesCFO Masterclass Series

---

*AI Workflows for the Modern CFO*

**PART 6**

# AI Board Reporting

*Use Case Two: From Close Data to Board Commentary*

Hindol Datta

*Chief Content and Product Officer, eFuturesCFO*

---

---

## A Note Before Part 6

---

Part 5 introduced the Finance Operations Copilot, the first production-grade AI workflow in John Campbell's deployment plan. The reader who completed Part 5 has now built a working version of that workflow and seen the eight-step methodology applied to a real workflow for the first time.

Part 6 introduces the second workflow: the AI Board Reporting and Narrative System. The use case is sequenced second for a specific reason that John articulated in his memo. The board reporting workflow builds on the close data that the Finance Operations Copilot has standardized, and Elena Vargas has already demonstrated proficiency with the underlying pattern through her informal use of ChatGPT for variance commentary. The team is ready for it. The data feeding it is clean. The governance precedent has been set in Part 5.

The reader will notice that Part 6 follows the same architectural pattern as Part 5: business problem, eight-step methodology, data walk-through, step-by-step tutorial, production discussion, expected outputs, assessment. The pattern is deliberate. By the end of the masterclass, having seen the pattern operate five times on five different workflows, the reader has internalized it.

What is different in Part 6 is the nature of the work. The Finance Operations Copilot in Part 5 worked at the transaction level, finding anomalies and patterns in granular data. The Board Reporting workflow in Part 6 works at the synthesis level, producing the narrative that an executive audience reads. The skills required are different. The Part 5 workflow demands analytical specificity. The Part 6 workflow demands editorial discipline.

The data package for Use Case Two, titled `Use_Case_2_Data.zip`, accompanies this PDF. It is even smaller than the Use Case One package because the data here is the close output rather than the transaction-level detail. The reader who has already worked through Part 5 will recognize the company, the cast, and the governance framework. The new files are the close package itself, the plan, the KPI history, the prior commentary, and the executive team's written notes for the upcoming board cycle.

*The Finance Operations Copilot finds the facts. The Board Reporting workflow tells the story. Both are essential. Neither substitutes for the other.*

Hindol Datta

# Contents

---

- Section 1** The Business Problem
- Section 2** The Eight-Step Methodology Applied
- Section 3** The Data Package Walk-Through
- Section 4** Step-by-Step Tutorial: Building the Workflow
- Section 5** From Tutorial to Production
- Section 6** Expected Outputs and Success Criteria
- Appendix A** Assessment: Twenty Questions
- Appendix B** Answer Key with Explanations

As with Part 5, Section 4 is the longest and most important. Plan to read this part with a computer nearby so you can follow along with the data files in either Claude or ChatGPT.

## Section 1 · The Business Problem

---

Every quarter, David Kim and Elena Vargas spend approximately one week constructing the board package. The package goes to the four directors of Helix Cloud Systems plus the audit committee chair, who reviews it with the same scrutiny that the directors of a public company would. The package runs approximately twenty-five to thirty slides plus exhibits. It is a serious document for a serious audience.

### What John heard in the discovery conversations

David described the construction process to John in his second week. The week begins on the Monday after close finalization. The team pulls the close numbers, computes the variances against plan, updates the KPI tracker, and produces the financial slides. They then begin the harder work: writing the CFO commentary.

The CFO commentary is the part of the board package that the audit committee chair reads most carefully. It is the part that communicates management judgment. Aaron Schultz, the outgoing interim CFO, did the commentary by writing each section from a blank document. He had no template that produced a first draft. Elena would help by drafting the variance explanations using ChatGPT informally, but the final structure and the executive judgment were always Aaron's alone.

The cycle compressed Elena and David in the two days before the board meeting. They worked late on Sunday and Monday. The package was usually finished early Tuesday morning, distributed Tuesday noon, and the meeting was Wednesday. The compression left no margin for revisions. If the board chair came back with a question on Monday evening, the team had to rush an answer.

### The cost of the current state

The visible cost is approximately one week of David's time per quarter, concentrated in the highest-leverage week of the period. It is also approximately four days of Elena's time. The work displaces other close-period responsibilities and creates the capacity ceiling pattern Karen Lindqvist flagged in Part 2.

The less visible cost is variability. Each board cycle produces commentary that reflects whoever drafted it on that particular cycle, with the editorial preferences and rhetorical habits of that drafter. A board chair reading commentary across cycles detects inconsistency. The inconsistency is not in the underlying judgment. It is in the prose. But it raises the question of whether the underlying judgment is also inconsistent.

The most consequential cost is that the commentary, in its current mode, is a backward-looking document. It explains what happened and why. It does not, in most cycles, surface forward-looking concerns with the discipline the audit committee chair has requested. The pattern is well known in finance functions: commentary that focuses on the period that closed is easier to write than commentary that examines what the period implies for periods ahead.

## What we will build

The AI Board Reporting workflow takes the close package data, the KPI history, the executive team notes, and the prior commentary as inputs, and produces a draft of the CFO commentary that David and the CFO will edit. The workflow does not produce the final commentary. It produces the structured first draft that the CFO edits down to the final version.

In its tutorial version, which you will build in Section 4, the workflow is a structured conversation with Claude or ChatGPT, anchored in the actual close data and the prior commentary. The reader will see the model produce a complete draft commentary in roughly the format that the prior CFO used, with the variances explained, the KPI movements interpreted, and the forward-looking concerns surfaced. The draft is not the final commentary. The CFO still edits it. But the gap between blank page and reviewable draft has been compressed from days to minutes.

### What this workflow is and is not

The workflow is a structured way of producing a first draft of the CFO commentary, using the close data and the prior commentary as inputs, with the CFO retaining full editorial authority over the final document. It is not an autonomous agent. It does not send anything to the board on its own. It compresses the time the CFO and the Controller spend on drafting, so that more of that time is available for editorial judgment and for the deeper analytical work that should inform the commentary.

## Section 2 · The Eight-Step Methodology Applied

---

This section walks through the eight-step methodology from Part 4, as applied to the AI Board Reporting workflow. The reader who has completed Part 5 will recognize the pattern. The application is different.

### Step One: Articulate the business problem

The business problem is Section 1 of this part. In one sentence: "The CFO commentary in the board package consumes approximately one week of senior finance leadership time per quarter, varies in quality and tone across cycles, and tends toward backward-looking description rather than forward-looking analysis."

The function that owns this problem is finance, specifically the CFO with support from David Kim and Elena Vargas. The audit committee chair (Diana Reyes-Okonkwo) is the principal consumer of the commentary. She has explicitly asked for forward-looking discipline and for consistency across cycles.

### Step Two: Specify the output

The workflow produces three categories of output.

First, the draft CFO commentary itself, in approximately the structure Aaron Schultz used: executive summary, revenue and bookings, gross margin, operating expenses, cash and runway, items for board attention, and management confidence. Format: flowing prose suitable for inclusion in the board package. Length: approximately three to five pages.

Second, a structured variance table covering every line item in the P&L; where actual differs from plan by more than five percent or by more than ten thousand dollars. Format: table with line item, plan, actual, variance dollar, variance percent, and a one-sentence explanation. Used as supporting exhibit and as fact-check for the prose commentary.

Third, a list of forward-looking concerns. These are items where the data suggests a future problem even if the current period looks acceptable. Format: numbered list with the concern, the supporting data, and the recommended action. The CAC payback extension is the example the reader will encounter.

### Step Three: Classify risk

The workflow is classified Tier Three. Its outputs are read by the board and influence directly the board's understanding of the company's performance. The outputs do not directly enter the financial statements, which keeps the workflow below Tier Four. But the consequences of an error in board commentary — misrepresentation, omission, lack of clarity — are material enough to warrant Tier Three controls.

Tier Three governance applies: complete audit trail, named human reviewer with finance team authority, monthly quality monitoring, CFO approval before deployment, audit committee notification, and review at every audit committee meeting.

## Step Four: Map data flow

The workflow reads, in production, the following data sources.

Input	Source	Classification
Close Package Actuals	NetSuite via MCP	Confidential
Close Package Plan	Planning system via MCP	Confidential
Prior Period Actuals	NetSuite via MCP	Confidential
KPI History	Looker via MCP	Confidential
KPI Definitions	Internal docs	Internal
Prior Board Commentary	Internal docs	Confidential
Board Package Outline	Internal docs	Internal
Executive Team Notes	Internal docs	Confidential

All inputs flow through approved environments under the AI governance framework. No data is sent to a public AI provider in a way that would result in training on the data.

## Step Five: Design the human review pattern

The workflow uses Review-before-output, the default pattern. The draft commentary, the variance table, and the forward-looking concerns list are all reviewed by named humans before any version of the document is shared beyond the finance team.

Named reviewers: John Campbell (primary), David Kim (secondary). The CFO has primary review authority because the commentary is his voice and his judgment. David has secondary review for factual accuracy against the underlying data. Elena Vargas manages the workflow operationally and prepares the draft for the CFO's review.

Review activities: verification (does the commentary accurately reflect the underlying data), judgment (does the commentary make the right calls about what to surface), and compliance (does the commentary meet the structural and tone standards of the prior commentary). Escalation: anything that cannot be resolved at the CFO review escalates to the audit committee chair before the meeting.

## **Step Six: Specify the audit trail**

Every workflow invocation produces the thirteen-field audit trail from Section 4 of Part 4. Retention is seven years. The audit trail captures the inputs, the prompts used, the draft output, the CFO's edits, and the final commentary as published. The auditor (Marsh and Henning) can reconstruct any board cycle on request.

## **Step Seven: Define the substitution path**

Primary model: Claude. Alternative: GPT. The workflow is designed to swap with approximately two engineering weeks of work. Annual testing of the substitution path applies, per framework.

## **Step Eight: Approval and registration**

Approval obtained from the Governance Working Group on day twelve of week eight of John's tenure. The workflow is registered as WFR-002 with version 1.0. The Q2 2026 board package will be the first cycle in which the workflow is used in pilot mode.

## Section 3 · The Data Package Walk-Through

---

The data package accompanying this PDF is titled `Use_Case_2_Data.zip`. Open it before continuing. Inside the zip you will find the common folder you already know from Use Case One, plus a new folder titled `uc2_board_reporting`. This section walks through what each file in the new folder contains.

### The eight Use Case Two files

#### **`close_package_actuals.csv`**

Approximately fifty rows showing the April 2026 P&L; plus key metrics. Each row has a line item, a category, the actual amount, and a note where one applies. This is the close output from the Finance Operations Copilot pipeline. It is the principal input to the board reporting workflow.

#### **`close_package_plan.csv`**

The same approximately fifty rows but with the plan numbers for April 2026. The reader will compare actuals to plan to identify variances.

#### **`prior_period_actuals.csv`**

Thirty rows covering both March 2026 (for sequential comparison) and April 2025 (for year-over-year comparison). The workflow uses this for trend context.

#### **`kpi_actuals.csv`**

Seventy-two rows covering twelve KPIs across six quarters (Q4 2024 through Q1 2026). This is the KPI history that underlies the board KPI table.

#### **`kpi_definitions.md`**

The methodology document defining how each of the twelve board KPIs is calculated. This is the reference the workflow uses when commenting on KPI movements. The document also contains two methodology notes that the workflow will surface as consistency flags.

#### **`prior_board_commentary.md`**

A sample of how the prior CFO (Aaron Schultz) wrote the Q4 2025 CFO commentary. This is the stylistic and structural reference for the workflow. The new CFO inherits this template, refines it, but starts from this base.

#### **`board_package_outline.md`**

The standard structure of the board package. Useful as context for the workflow to understand where the CFO commentary fits in the overall package.

## **executive\_team\_notes.md**

Brief written inputs from each function head ahead of the board package. These are the qualitative additions to the quantitative data, providing context on what is going on in each function. The workflow incorporates these into the draft commentary.

## **How to look at the files**

Before continuing to the tutorial, spend approximately ten minutes reviewing the files.

Open `close_package_actuals.csv` in a spreadsheet. Notice the structure: revenue lines summing to Total Revenue, COGS lines summing to Total COGS, operating expense lines grouped by category (S&M;, R&D;, G&A;). Read down the column. Notice the note on the Legal Fees line: "NOTE: elevated this month, see commentary." This is the kind of human annotation the workflow will pick up.

Open `close_package_plan.csv`. Compare any line you remember from actuals. Some are close (revenue is slightly above plan). Some are significantly different (legal fees are an order of magnitude above plan).

Open `kpi_actuals.csv`. Look at the CAC Payback rows. Notice the trajectory across the six quarters: 14.2 to 15.1 to 16.4 to 17.8 to 18.5 to 19.4. Five quarters of extension. This is a pattern that a workflow scanning the KPI history should detect.

Open `prior_board_commentary.md`. Read the first two sections (the executive summary and the revenue paragraph). Notice the voice: measured, declarative, structured. This is the voice the workflow should produce.

Open `executive_team_notes.md`. Skim through the notes from each function head. Notice the variety: some are quantitative, some are qualitative, some surface concerns, some surface positive developments. The workflow will need to synthesize these into the CFO commentary.

### **Reading the prior commentary carefully**

Of all the files in this package, the prior board commentary is the most important for the workflow. The commentary is both the template and the voice. The workflow will produce a draft that matches the structure and the tone of the prior commentary, and your editorial review will compare the draft against the template to assess fidelity.

## Section 4 · Step-by-Step Tutorial: Building the Workflow

This is the heart of Part 6. You will build a working version of the AI Board Reporting workflow using either Claude or ChatGPT. The tutorial works with either tool. The pattern is the same pattern you learned in Part 5: establish the role, load the data with verification, ask specific analytical questions with format requirements, synthesize, and document.

Before you begin, have the following ready: a current web browser, an account on either Claude (claude.ai) or ChatGPT (chatgpt.com), and the Use\_Case\_2\_Data zip downloaded and unzipped. Plan for approximately fifty minutes to complete the tutorial. The tutorial has eleven steps. The middle steps are longer than in Part 5 because the analytical questions are more complex.



### Tutorial Step 1: Start a fresh conversation

Open your web browser. Navigate to claude.ai or chatgpt.com. Sign in. Start a new conversation. Do not continue the conversation from Part 5. Each workflow operates in its own conversation so that the model is not confused by context from a different workflow.

#### *Why this matters:*

*Workflows are independent. The Finance Operations Copilot from Part 5 and the Board Reporting workflow you are about to build have different purposes, different inputs, different audiences, and different output formats. Mixing them in the same conversation produces inferior output for both. The discipline of one workflow per conversation is the same discipline that governs the workflow registry in production.*

### Tutorial Step 2: Set the role and context

In the input box, paste the following prompt. Read it before sending.

```
You are acting as my AI Board Reporting Workflow for a Series B B2B SaaS company called Helix Cloud Systems. I am the CFO, John Campbell. The company has $22.4M in annual recurring revenue, 142 employees, and reports quarterly to a board of four directors plus an audit committee chair.
```

```
Your job is to help me produce the CFO commentary section of the board package. The commentary is the most editorially demanding part of the package. It is read with particular attention by the audit committee chair, Diana Reyes-Okonkwo, who has explicitly asked for forward-looking discipline and for consistency in tone and structure across cycles.
```

You will operate in three phases:  
 Phase 1: Receive the data files. I will paste them one at a time.  
 Phase 2: Conduct analysis. I will ask you specific questions in sequence.  
 Phase 3: Produce the draft commentary. I will give you the format requirements.

You operate under a governance framework that requires every output to be reviewed by a named human (me) before any version is shared beyond the finance team. The commentary you produce is always a draft. I will edit it before publication.

Please confirm you understand the role and the three-phase structure, and then ask me to paste the first data file.

Press send.

### **Why this matters:**

*The role prompt is more detailed in this workflow than it was in Part 5, because the editorial requirements are more specific. The model needs to understand not just what kind of analysis to do but what kind of writing to produce. The mention of Diana Reyes-Okonkwo and her expectations is deliberate. The model uses this information to shape the tone of the output. Tell the model who the audience is, and the model writes for that audience.*

Expected response: the model confirms it understands the role and asks for the first file. Some models will summarize back their understanding, which is fine. Read the response. If anything is misunderstood, correct it before proceeding.

## **Tutorial Step 3: Provide the board package outline**

Open board\_package\_outline.md from the uc2\_board\_reporting folder. Copy the entire contents. Paste into the conversation with this prompt.

Here is the standard structure of our board package. Please read it carefully and confirm you understand which section of the package you will be producing (Section 7, the CFO Commentary). After reading, tell me:

1. What is the audience for the package?
2. What section will you be producing?
3. How long should that section be?

[Paste the contents of board\_package\_outline.md here]

### **Why this matters:**

*You are giving the model the structural context of the document it will contribute to. The CFO commentary is one section of a larger package. Understanding the place of that section in the whole shapes how the model writes it. The questions at the end verify that the model has read the document carefully rather than skimmed it.*

Expected response: the model confirms it will be producing the CFO commentary (Section 7), at approximately three to five pages in length, for an audience of the board of directors plus the audit committee

chair.

### Tutorial Step 4: Provide the prior board commentary

Open `prior_board_commentary.md`. Copy the entire contents. Paste into the conversation with this prompt.

Here is the Q4 2025 board commentary, written by the prior interim CFO, Aaron Schultz. This is the stylistic and structural reference for what you will produce. Please read it carefully and tell me:

1. What is the structural pattern of the commentary (how many sections, what does each cover)?
2. What is the voice and tone?
3. What rhetorical habits do you observe (for example, how does the writer handle uncertainty, how does the writer handle bad news)?
4. What should be different about the commentary I write now, given that I am a new CFO with different priorities (forward-looking discipline, attention to consistency, my Series C window)?

Be specific. Quote phrases from the prior commentary where useful.

[Paste the contents of `prior_board_commentary.md` here]

#### *Why this matters:*

*This step is crucial. The model is not just being asked to imitate the prior commentary. The model is being asked to understand it, articulate its patterns, and identify where the new CFO's priorities differ. This forces the model to think about the commentary as a piece of executive writing with a specific purpose, not as a template to be filled in. The fourth question is especially important: it acknowledges that you are a different CFO with different emphasis, and the commentary you produce will reflect that difference.*

Expected response: the model identifies approximately seven sections in the prior commentary (executive summary, revenue and bookings, gross margin, operating expenses, cash and runway, items for board attention, management confidence). It characterizes the voice as measured, declarative, and willing to acknowledge uncertainty. It identifies habits like the use of phrases such as "we do not yet have a confident explanation" when management does not know something. For the fourth question, it identifies the new CFO's emphasis on forward-looking discipline, on consistency, and on the Series C window.

### Tutorial Step 5: Provide the close package actuals and plan

Open both `close_package_actuals.csv` and `close_package_plan.csv`. You will paste them in a single prompt for efficiency. Copy the actuals file contents first, then in your prompt include both.

Here are the April 2026 close numbers. The first table is the actuals; the second is the plan. Each table has line items, categories, and amounts.

Please confirm receipt and produce a variance analysis:

1. Compute the variance of actual to plan for each line item (in dollars and percent)
2. Identify all line items where the absolute variance exceeds 5% or \$10,000

3. Group the material variances by category (Revenue, COGS, S&M, R&D, G&A)
4. For each material variance, do not yet explain it – just identify it

I will provide the explanations from the executive team notes in the next step.

ACTUALS:

[Paste contents of close\_package\_actuals.csv]

PLAN:

[Paste contents of close\_package\_plan.csv]

### **Why this matters:**

*You are now asking the model to do the quantitative work that would otherwise take Elena Vargas several hours: computing variances, identifying material ones, and structuring them by category. Notice the explicit instruction "do not yet explain them." This forces the model to first lay out the facts before attempting interpretation. In drafting work, the discipline of separating facts from interpretation is what produces commentary that survives editorial review.*

Expected response: the model produces a variance table showing the material variances. The big ones to verify: Total S&M; at +\$80K versus plan (+13.6%); Total R&D; at -\$47K versus plan (-7.8%); Legal Fees at +\$114K versus plan (+950%); Total G&A; at +\$123K versus plan (+29.3%); Total Revenue slightly favorable. Read the response carefully and verify the numbers match what you would compute yourself.

## **Tutorial Step 6: Provide the executive team notes**

Open `executive_team_notes.md`. Copy the entire contents. Paste into the conversation with this prompt.

Here are the written inputs from each function head ahead of the May 2026 board package. Please:

1. Read each note carefully
2. For each material variance you identified in the prior step, find any explanation in the function head notes
3. Identify which variances are not yet explained in the notes (these are gaps I need to follow up on before publication)
4. Identify any items in the notes that are not yet reflected in the close numbers but that the board should know about

Do not yet write the commentary. We are still in the analysis phase.

[Paste contents of executive\_team\_notes.md here]

### **Why this matters:**

*You are now bringing in the qualitative context. The function heads have written narrative notes about what is happening in their areas. The model's job is to match these notes to the quantitative variances. This is the cross-referencing work that takes Elena Vargas the longest in current practice. The model does it in seconds. Notice that the prompt asks for both the explanations that exist (where notes match variances) and the gaps*

(where notes do not explain variances). The gaps are where you need to follow up before publication.

Expected response: the model matches notes to variances. Examples of expected matches: Naomi's note explains the legal fees overspend (Series C prep and customer contract negotiation); Raj's note explains the R&D; underspend (delayed engineering hires); Michael's note partly explains the sales commissions overspend (April bookings ahead of plan) but does not directly explain the advertising overspend. The advertising overspend is one of the gaps the model should flag.



## Tutorial Step 7: Provide the KPI history and definitions

Open both `kpi_actuals.csv` and `kpi_definitions.md`. Paste them in one prompt.

```
Here is the KPI history (six quarters of twelve KPIs) and the KPI definitions document.
```

```
Please:
```

1. Read the definitions document carefully
2. Identify any methodology notes or consistency concerns flagged in the document
3. For each of the twelve KPIs, characterize the trajectory across the six quarters
4. Identify any KPIs that should be flagged for the board because their trajectory is concerning
5. Identify any KPIs where you need to clarify methodology before publication

```
KPI HISTORY:
```

```
[Paste contents of kpi_actuals.csv]
```

```
KPI DEFINITIONS:
```

```
[Paste contents of kpi_definitions.md]
```

### *Why this matters:*

*This is the largest data ingestion of the workflow so far. The KPI history is the trend data. The definitions document is the methodology. The two together let the model interpret the KPI movements correctly.*

*Question 2 is particularly important: the definitions document contains explicit methodology notes that the workflow should surface. One of them, the gross margin methodology inconsistency between Q4 2025 and Q1 2026, is exactly the kind of consistency issue Diana Reyes-Okonkwo will flag if you do not flag it first.*

Expected response: the model identifies the NDR reported-versus-expansion-only restatement note, and the gross margin methodology inconsistency between Q4 2025 and Q1 2026. For the KPI trajectories, the model should identify ARR growing steadily, NDR holding at 118%, gross margin in the 77-79% range, but critically, CAC Payback extending from 14.2 months to 19.4 months across the six quarters. The model should flag CAC payback as a concerning trajectory.

## Tutorial Step 8: Provide the prior period actuals

Open `prior_period_actuals.csv`. Paste with this prompt.

Here are the prior period actuals: March 2026 (the immediately preceding month) and April 2025 (the same month one year prior). Please:

1. For each material variance you identified versus plan, compute the variance versus the prior month and versus the prior year
2. Identify any variances that are large versus plan but small versus the prior period (which suggests plan was wrong)
3. Identify any variances that are large versus plan AND versus the prior period (which suggests actual is unusual)
4. Identify any items where the prior year comparison is materially different from the sequential month comparison

[Paste contents of `prior_period_actuals.csv`]

### Why this matters:

*Variance commentary is more credible when it is multi-dimensional. A variance against plan tells you that the period differed from expectations. A variance against the prior month tells you the direction the business is moving. A variance against the prior year tells you the longer-term trajectory. The model now has all three. Notice question 2: a variance that is large versus plan but small versus prior period often means the plan was set too aggressively. This distinction matters when writing commentary, because the right framing depends on the cause.*

Expected response: the model produces a multi-dimensional variance view. The legal fees overspend should now show as significant versus plan, versus March (where legal was already elevated), and versus April 2025 (where legal was at baseline). This three-way comparison tells the full story of the legal fees climb that the reader first saw in Part 5.

## Tutorial Step 9: Produce the variance commentary

Now you ask the model to produce the first piece of the actual commentary.

Now I want you to draft the operating expense commentary section of the board commentary. Use the prior board commentary's voice and structure as your reference, but adapt the content to the April 2026 reality.

The section should cover:

1. S&M overspend and what drove it (variance is 13.6% over plan)
2. R&D underspend and why it should not be read as under-investment (variance is 7.8% under plan)
3. G&A overspend driven primarily by legal fees (variance is 29.3% over plan, with legal driving most of the gap)

Style requirements:

- Three to five paragraphs of measured prose, no bullet points
- Each variance is named, then explained with the function head's reasoning, then contextualized

- Where Naomi's note about legal fees applies, incorporate it without quoting it verbatim
- Where you do not have a confident explanation (such as the advertising overspend), acknowledge that openly
- The voice should match Aaron Schultz's voice from the prior commentary
- No defensive language, no marketing language, no exclamation points

**Why this matters:**

*This is the first piece of actual writing the workflow produces. Notice the precision of the requirements. The model is told what content to cover, what style to use, what to acknowledge about uncertainty, and what to avoid. The result is a draft that requires editorial review but is far closer to publishable than anything you would produce from a blank page in the same amount of time. Notice also the instruction to "incorporate without quoting verbatim" Naomi's note. The model is being asked to translate function-head notes into CFO commentary, which is the editorial work the new CFO would otherwise have to do himself.*

Expected response: a three-to-five paragraph commentary that reads like the prior commentary in voice and structure but covers the April 2026 content. Read it carefully. Edit anything that feels off. The draft is yours to revise. You are the CFO. The commentary is your voice.

**Tutorial Step 10: Produce the forward-looking concerns section**

Now you ask the model to draft the section the audit committee chair cares about most: the forward-looking concerns.

Now draft the "Items for Board Attention" section. This is the section that the audit committee chair, Diana Reyes-Okonkwo, reads most carefully. It is also the section where my emphasis differs from Aaron's: I want it to be forward-looking, surfacing concerns that the data suggests will become problems in future periods, even if they are not problems in the current period.

For this section, address at least the following items:

1. The CAC payback extension trend (the five-quarter trajectory in the KPI data)
2. The Magic Number trend (also visible in KPI data)
3. The Falcon Group collections risk (from Michael's note)
4. The Beacon Capital yellow health flag (from Priya's note)
5. The Series C window timing relative to forecast predictability
6. The methodology inconsistency on Gross Margin (Q4 2025 vs Q1 2026)

For each item:

- State the item in one sentence
- Explain what the data shows
- Acknowledge what we do not yet know
- Recommend what the board should expect to hear from management in the next cycle

Style: measured executive prose. No bullets within the discussion of each item. Each item gets one paragraph. The section opens with one sentence framing why these items matter.

**Why this matters:**

*This is the most editorially demanding step in the workflow. The model is being asked to surface concerns that may not be obvious from a quick read of the close numbers, to acknowledge uncertainty honestly, and to recommend specific follow-up. These are executive skills, not technical skills. The model can do them when given clear guidance, which is what the prompt provides. The six items are all things you would surface yourself if you had the time. The model has compressed the analysis into seconds. Your job is to verify and refine.*

Expected response: a structured forward-looking concerns section covering the six items. The CAC payback discussion is the most consequential. The model should clearly identify the five-quarter extension, acknowledge that the cause is not fully understood, and recommend that management return to the board in the next cycle with analysis. The methodology inconsistency item is the kind of self-flag that builds credibility with the audit committee chair.

**Tutorial Step 11: Synthesize and document**

You now ask for a final executive summary and close the loop.

Now produce the executive summary, which appears at the top of the commentary even though we wrote it last. The executive summary is read carefully by all board members and is the only section some of them read in detail.

**Requirements:**

- Three to five paragraphs of prose, no bullets
- Open with one sentence that frames the period
- Cover the period's headline performance (revenue, ARR, key KPIs)
- Identify the two or three items where my attention is most directed
- Identify the one item where my confidence is lowest
- Close with one sentence that frames what the rest of the commentary will cover

Use the same voice and tone we have been developing through this workflow.

After you receive the executive summary, do two things.

First, take all four pieces (executive summary, operating expense commentary, items for board attention, and the KPI variance analysis) and assemble them in the order they will appear in the final document. The order is: executive summary, then variance analysis exhibits, then operating expense commentary, then items for board attention. You may want to ask the model to produce the assembled document with appropriate transitions.

Second, save the conversation and write a brief note for the workflow registry. Note the date, the data files used, the prompts that worked best, and any places where the model required iteration. As with Part 5, the documentation discipline matters as much as the workflow itself.

**Why this matters:**

*The executive summary is written last for a reason. You cannot know what to include in the summary until you have done the underlying analysis. The pattern of writing the summary last and placing it first is what*

*good executive writers do. By specifying this pattern explicitly in the workflow, you ensure the model produces a summary that reflects the actual analysis rather than a summary that pretends to.*



## What you have just done

You have built the second AI workflow of the masterclass. Together with the Finance Operations Copilot from Part 5, these two workflows now cover the two largest time consumers in David and Elena's quarterly cycle: the close transaction review (Part 5) and the board commentary drafting (Part 6).

The pattern you used in Part 6 is similar to the pattern in Part 5 but with more emphasis on editorial discipline than on analytical specificity. You established the role with more context about the audience. You loaded the data in phases, starting with the structural context (the package outline), then the stylistic reference (prior commentary), then the quantitative data (actuals, plan, prior periods, KPIs), then the qualitative context (executive team notes). You asked the model to write in sections (operating expense commentary, forward-looking concerns, executive summary) rather than all at once.

In production, this workflow runs in roughly two hours every quarter, compared to roughly five days currently. The CFO's editorial time on the commentary becomes the substance of the work, not the procedural overhead.

*The model writes the draft. The CFO writes the final. The compression is in the drafting. The editorial judgment remains where it belongs.*

## Section 5 · From Tutorial to Production

---

The tutorial version of the AI Board Reporting workflow is, like Part 5's copilot, the right place to start. The production version of this workflow is a more significant engineering investment than the copilot, because the data ingestion is more varied and the output structure is more elaborate.

### What changes in production

#### 1. The data flows automatically from multiple sources

In the tutorial, you copied data from CSV files and markdown files and pasted into the chat. In production, the workflow pulls the close numbers from NetSuite via MCP, the KPI data from Looker via MCP, the executive team notes from a structured intake form in your internal collaboration tool, and the prior commentary from your document repository. All inputs flow without copy and paste.

#### 2. The workflow runs on the board package schedule

The workflow is invoked automatically on day five of close, after the close numbers are finalized and after the function heads have submitted their notes. The draft commentary is in the CFO's review queue by mid-morning on day five, well before the compression that currently happens in the two days before the board meeting.

#### 3. The audit trail is automatic and integrated with the package

In the tutorial, you wrote a manual note for the workflow registry. In production, every invocation is logged with the thirteen-field audit trail. The audit trail is associated with the specific board cycle, so the auditor can reconstruct the commentary drafting process for any quarter on request.

#### 4. The outputs flow into the board package directly

In the tutorial, the commentary is prose in the chat. In production, the commentary is delivered as a draft document in the format your board package expects (a Google Doc or a Word document with the appropriate structure), already populated with the variance table, the KPI table, and the prose commentary.

### Engineering work to get from tutorial to production

Roughly eight to ten engineering weeks for the initial production version. The work is somewhat larger than the Part 5 copilot because of the additional data sources and the more elaborate output format.

Three weeks for the data integration: NetSuite via MCP, Looker via MCP, the document repository, and the executive team note intake form. Two weeks for the workflow orchestration. Two weeks for the output document generation, including the formatting that matches the board package template. One week for governance integration and the audit trail wiring. Two weeks for testing, calibration, and iteration through at

least one full board cycle.

## What you can do before the production version is ready

As with the Finance Operations Copilot, you can run the tutorial version every quarter while the production version is being built. The first quarter you run the tutorial version, plan to spend the full day on it: roughly two hours running the workflow in the chat, then a full afternoon editing the draft commentary into final form. The second time, the work compresses meaningfully because you know the prompts and you know what to expect.

John Campbell's plan is to use the tutorial version for the Q2 2026 board package (the first cycle in which he is the CFO), with Elena Vargas operating the workflow and David Kim and John reviewing the output. By the Q3 2026 cycle, the production version should be operational. By the Q4 2026 cycle, the workflow is the standard process.

### A note about voice

The voice of the AI-drafted commentary will not be exactly the voice of the CFO. Your editorial pass will adjust the voice toward your own. Over several cycles, you can refine the prompts to make the model produce drafts that are closer to your voice from the start. The discipline of continuous prompt refinement is part of running the workflow well in production.

## Section 6 · Expected Outputs and Success Criteria

The AI Board Reporting workflow produces specific deliverables on a quarterly cadence. The success criteria are measured against both efficiency and quality.

### Expected outputs

Every board cycle, the workflow produces four artifacts.

#### Draft executive summary

A three-to-five paragraph executive summary covering the period's headline performance, the items for board attention, and the framing for the rest of the commentary.

#### Draft CFO commentary (full)

The complete commentary covering revenue and bookings, gross margin, operating expenses, cash and runway, items for board attention, and management confidence. Three to five pages of measured executive prose.

#### Structured variance analysis

A table covering every line item where actual differs from plan by more than five percent or ten thousand dollars, with the plan, the actual, the variance dollar and percent, and a brief explanation of the driver where the executive team notes provide one. Used as supporting exhibit and as fact-check for the commentary.

#### Forward-looking concerns list

A numbered list of items where the data suggests a future concern, even if the current period is acceptable. The CAC payback extension is the example the reader saw in this part. Each item has supporting data and a recommended follow-up.

### Success criteria

Metric	Baseline	Target (12 months)
Commentary drafting time	4-5 days per quarter	1 day per quarter
Cycle predictability	Compression in final 48 hours	Steady through cycle
Editorial revisions per cycle	40-60% of draft rewritten	15-25% of draft rewritten
Audit committee chair feedback	Specific items per cycle	Decreasing item count

Metric	Baseline	Target (12 months)
Forward-looking concerns surfaced	1-2 per cycle (sometimes)	3-5 per cycle (consistently)

## What is being measured, and why

Each metric reflects a specific aspect of the value the workflow produces.

Commentary drafting time is the principal efficiency measure. The compression from days to roughly one day is the headline benefit.

Cycle predictability is the qualitative measure of stress. The current cycle is stressful because everything compresses into the two days before the board meeting. The new cycle should distribute the work evenly across the full week.

Editorial revisions per cycle measure draft quality. The first few cycles will require heavy revision (40-50%). As the prompts are refined and the model learns your editorial preferences, the revision percentage should decline.

Audit committee chair feedback is the external quality measure. Diana's feedback in each cycle gives you direct signal on what the audit committee is responding to. Decreasing feedback over time signals that the commentary is meeting expectations.

Forward-looking concerns surfaced is the most important measure because it captures the structural improvement in the commentary, not just the efficiency gain. A commentary that consistently surfaces three to five forward-looking concerns per cycle is a commentary that meets Diana's explicit request.



# End of Part 6

## *AI Board Reporting and Narrative System*

You have built the second AI workflow of the masterclass. Together with Part 5, you now have two production-capable workflows addressing the two largest time-consuming activities in the finance cycle: transaction review during close, and board commentary drafting at quarter end.

In Part 7, the masterclass turns to the third use case: Pipeline Intelligence. This workflow operates outside the finance function in the sense that Michael O'Brien and the sales operations function own the deployment. But the workflow consumes pipeline data and produces forecasts that feed directly into the finance function's revenue forecasting. By the end of Part 7, you will have built a third workflow that connects the sales operations function with the finance function, with the same governance discipline you have learned.

Before proceeding, take the assessment that follows. The questions test both the substance of the Board Reporting workflow and the editorial discipline it represents.



# Appendix A · Assessment

---

Twenty questions on Part 6. Twelve multiple choice, five short answer, three scenario-based. Answer key follows.

## Part I: Multiple Choice

**1. The AI Board Reporting workflow is sequenced second in John Campbell's deployment plan because:**

- (a) It is the most time-consuming workflow.
- (b) It builds on close data the Finance Operations Copilot has standardized, and Elena Vargas has informal proficiency with the underlying pattern.
- (c) The board specifically requested it.
- (d) It is the lowest-risk workflow.

**2. The Board Reporting workflow is classified as:**

- (a) Tier One: Internal Exploratory.
- (b) Tier Two: Internal Operational.
- (c) Tier Three: Financial Reporting Adjacent.
- (d) Tier Four: Financial Reporting Direct.

**3. The primary reviewer for the commentary draft is:**

- (a) Elena Vargas.
- (b) David Kim.
- (c) John Campbell.
- (d) Sarah Chen.

**4. The Q4 2025 board commentary was written by:**

- (a) John Campbell.
- (b) Sarah Chen.
- (c) Aaron Schultz, the prior interim CFO.
- (d) Elena Vargas as a draft.

**5. The seeded S&M; variance is:**

- (a) S&M; came in 3% under plan, a favorable variance.
- (b) S&M; came in 13.6% over plan, an unfavorable variance driven primarily by Advertising and Events.
- (c) S&M; was exactly at plan.
- (d) S&M; came in 25% over plan due to commissions.

**6. The seeded R&D; variance is:**

- (a) R&D; came in 7.8% under plan, driven by delayed engineering hires.
- (b) R&D; came in 12% over plan, driven by contractor spend.
- (c) R&D; came in exactly at plan.
- (d) R&D; came in 20% under plan, suggesting under-investment.

**7. The CAC payback trajectory over the six quarters of KPI history shows:**

- (a) Steady at approximately 15 months across all six quarters.
- (b) Decreasing from 19 months to 14 months over the period.
- (c) Extending from 14.2 months to 19.4 months over the period.
- (d) Unstable with no clear pattern.

**8. The methodology inconsistency flagged in the KPI definitions document concerns:**

- (a) ARR calculation between Q3 and Q4 2025.
- (b) NDR reporting basis (with vs. without escalators).
- (c) Gross Margin treatment of Customer Success salaries between Q4 2025 and Q1 2026.
- (d) Magic Number calculation methodology.

**9. The "executive team notes" file contains:**

- (a) Detailed financial statements from each function.
- (b) Brief written inputs from each function head ahead of the board package.
- (c) Verbatim minutes from the executive team meeting.
- (d) A list of action items from the prior board meeting.

**10. The forward-looking concerns section of the commentary should:**

- (a) Restate the period's headline performance.
- (b) Surface items where the data suggests a future concern, even if the current period is acceptable.
- (c) List the action items from the prior board meeting.
- (d) Provide a forecast for the next four quarters.

**11. In production, the workflow runs:**

- (a) Automatically on day five of close.
- (b) When the CFO manually invokes it.
- (c) On a daily schedule.
- (d) Only at year-end.

**12. The target reduction in commentary drafting time is approximately:**

- (a) From 4-5 days per quarter to 2 days per quarter.
- (b) From 4-5 days per quarter to 1 day per quarter.
- (c) From 4-5 days per quarter to zero (full automation).
- (d) No specific target has been set.

**Part II: Short Answer**

13. In two or three sentences, explain why the tutorial loads the prior board commentary file before loading any quantitative data. What does the order of file loading accomplish?

14. The tutorial instructs the model to "identify the variances but do not yet explain them" in Step 5, and then provides the explanations from executive team notes in Step 6. In two or three sentences, explain why separating facts from interpretation matters in board commentary drafting.

15. The methodology inconsistency on Gross Margin between Q4 2025 and Q1 2026 is a small finding that the workflow flags. In two or three sentences, explain why such small inconsistencies are worth surfacing in board commentary even when the consequences are minor.

16. The Part 6 workflow has the same Tier Three classification as the Part 5 workflow, but the human review pattern places the CFO (not the Controller) as primary reviewer. In two or three sentences, explain why this workflow places the CFO in primary review.

17. The CAC payback extension is identified as a "forward-looking concern" rather than a "current period variance." In two or three sentences, explain the distinction and why it matters for how the item is presented to the board.

### Part III: Scenario-Based

18. Scenario: At the Q2 2026 board meeting, after the first cycle in which you used the AI Board Reporting workflow, Diana Reyes-Okonkwo pulls you aside. She tells you the commentary was the cleanest she has seen in three quarters, and asks whether the model could be given a more aggressive role going forward — perhaps drafting the entire board package, not just the CFO commentary. In one paragraph of executive prose, describe how you would respond, what governance principle you would invoke, and what you might offer as a constructive next step.

19. Scenario: Three quarters into using the workflow in production, Elena Vargas notices that the workflow consistently misses one type of finding. Customer churn losses that are unanticipated relative to the health scoring system never appear in the forward-looking concerns section, because the workflow has no access to the customer health scoring data. Priya Sundaram's notes mention churn at a high level, but the underlying patterns are invisible to the workflow. In one paragraph, describe how you would address this gap, including any changes to the workflow design and any governance steps you would take before deploying the changes.

20. Scenario: The audit committee, after reviewing your three-quarter track record with AI-drafted board commentary, asks whether the practice should be disclosed in the company's Series C diligence materials. The institutional investor leading the Series C round is known to be sophisticated on AI matters and may either welcome or scrutinize the disclosure. In one paragraph of executive prose, describe what you would recommend to the audit committee, what you would include in the disclosure if you recommend disclosing, and why.

## Appendix B · Answer Key with Explanations

---

Multiple choice answers are noted. Short answer and scenario discussions follow.

### Multiple Choice Answers

**Question 1: (b)**

The workflow is sequenced second because it builds on the close data the Finance Operations Copilot has standardized in Q1, and because Elena Vargas has demonstrated proficiency with the underlying pattern through her informal ChatGPT use for variance commentary. See Section 8 of Part 3 and Section 1 of Part 6.

**Question 2: (c)**

Tier Three because the commentary is read by the board and the audit committee chair, influencing their understanding of the company's performance. The outputs do not directly enter the financial statements, which keeps it below Tier Four.

**Question 3: (c)**

John Campbell is primary reviewer because the commentary is his voice and his judgment. David Kim is secondary reviewer for factual accuracy. Elena Vargas operates the workflow and prepares the draft for review.

**Question 4: (c)**

The Q4 2025 commentary was written by Aaron Schultz, the prior interim CFO, before John Campbell joined in April 2026. The commentary is used as a stylistic and structural reference.

**Question 5: (b)**

S&M; came in \$80,000 over plan, a 13.6% unfavorable variance. The principal drivers are Advertising (\$23K over plan, +41.8%) and Events (\$14K over plan, +77.8%), with smaller overspends in Sales Salaries and Sales Commissions.

**Question 6: (a)**

R&D; came in \$47,000 under plan, a 7.8% favorable variance. The driver is two engineering hires that delayed to May, as noted by Raj Patel in the executive team notes.

**Question 7: (c)**

CAC payback extended from 14.2 months in Q4 2024 to 19.4 months in Q1 2026, a 5-quarter linear trajectory. This is the leading-indicator concern the workflow should surface for the board.

**Question 8: (c)**

The Q4 2025 board package calculated Gross Margin excluding Customer Success salaries from COGS, while the Q1 2026 package included them. The two methodologies produce approximately a 4-5 percentage point spread. The new methodology is now standard but the inconsistency should be disclosed.

**Question 9: (b)**

The executive team notes are brief written inputs from each function head submitted to the CFO ahead of the board package. They are the qualitative context that complements the quantitative close data.

**Question 10: (b)**

Forward-looking concerns surface items where the data suggests a future problem, even if the current period is acceptable. The CAC payback extension is the example: each quarter individually was acceptable, but the trajectory is concerning.

**Question 11: (a)**

In production, the workflow runs automatically on day five of close, after the close numbers are finalized and after function heads have submitted their notes. The draft is in the CFO's review queue mid-morning of day five.

**Question 12: (b)**

The target is to reduce commentary drafting time from approximately 4-5 days per quarter to approximately 1 day per quarter. The compression comes from the workflow producing a structured first draft that the CFO edits, rather than drafting from a blank page.

## Short Answer Explanations

**13. Why load the prior commentary first**

The order is deliberate. Loading the prior commentary first establishes for the model what the output will look like before it sees any of the data that will populate the output. The model writes far better when it understands the target document at the start. Loading the data first would produce analysis that is structurally disconnected from the editorial output the workflow ultimately needs to produce. The deeper executive principle is that good drafting always starts with the document you are producing, not with the data you have available.

**14. Separating facts from interpretation**

Variance commentary is more credible when the facts are clearly distinguished from the explanations. A reader who cannot tell where the data ends and the interpretation begins cannot evaluate either. By forcing the model to first identify variances without explaining them, then layer in the function-head explanations in a separate step, the workflow produces commentary in which the data and the interpretation are clearly separated. This discipline also surfaces the gaps: variances for which no explanation exists in the notes are flagged for follow-up, rather than being papered over with plausible-sounding speculation.

### 15. Why small inconsistencies matter

Audit committee chairs are particularly sensitive to consistency issues because inconsistency in metric definitions can mask substantive changes in performance. The Gross Margin methodology shift between Q4 2025 and Q1 2026 is small in numerical terms but signals a methodology that is not yet locked down. Surfacing such issues proactively builds credibility with the audit committee; discovering them after the audit committee has flagged them undermines credibility. The deeper principle is that the value of self-flagging is asymmetric: it costs little when nothing is wrong, and it protects substantially when something is wrong.

### 16. Why the CFO is primary reviewer

The board commentary is the CFO's editorial voice on behalf of management. The Controller can verify factual accuracy, but only the CFO can decide whether the commentary expresses the right judgment in the right tone. Placing the CFO as primary reviewer ensures that the editorial authority sits with the executive whose name will appear at the bottom of the commentary. Placing the Controller as primary would be functionally appropriate but would create ambiguity about whose judgment the board is reading.

### 17. Variances versus forward-looking concerns

A variance is a measurable difference between actual and plan in the current period. A forward-looking concern is a pattern in the data that suggests a future problem, regardless of whether any current variance has flagged it. The CAC payback extension is the example: each individual quarter's CAC payback was acceptable, so it would not appear in a variance commentary. The trajectory across five quarters is the concern, and only a workflow that examines trends explicitly will surface it. The distinction matters because variance commentary alone tends to be backward-looking, and the audit committee chair has specifically requested forward-looking discipline.

## Scenario Discussions

### 18. Diana's expansion request

The right response is to thank Diana for the feedback, acknowledge the constructive intent, but decline to expand the workflow at this stage. The governance principle to invoke is incrementalism: every AI workflow deployment is approved against a specific scope, and expanding the scope without going back through the eight-step methodology would undermine the discipline that produced the quality result Diana is responding to. Specifically, drafting the entire board package would mean producing not just commentary but quantitative analysis, exhibit construction, and section assembly. Each of these requires its own design, its own review pattern, and its own audit trail. As a constructive next step, propose that we identify the two or three specific sections of the package where the current manual process is most time-consuming and least judgment-intensive, evaluate each against the framework, and bring proposals back to the audit committee for any expansion. The general principle is that early success creates pressure for over-extension, and the discipline of resisting that pressure is what makes the program durable over the eighteen-month horizon.

### 19. The customer churn blind spot

The right approach is to treat this finding as design refinement rather than as a flaw. The workflow was designed against the data sources available, and a missing data source is a fixable gap, not a workflow failure. Specifically, three things should happen. First, document the gap formally in the workflow registry: the workflow does not have access to customer health scoring data and therefore cannot surface health-driven churn risk patterns. This protects the workflow against the suggestion that it should have caught the patterns when it lacked the data to do so. Second, propose adding customer health data as an input to the workflow, with the engineering work scoped against the existing MCP integration to the customer success platform. Third, treat the change as a material modification under the governance framework: it requires re-review by the Governance Working Group, updates to the data flow map, and a re-test of the substitution path. The change is small in engineering terms but real in governance terms. Until the change is approved and deployed, document the limitation in every board cycle's workflow output so that the CFO and the audit committee know what the workflow does and does not see.

### 20. Series C disclosure of AI use

The recommendation should be to disclose the practice deliberately, in a section of the diligence materials specifically addressing the AI governance program. A sophisticated institutional investor will either ask the question or notice the absence of the topic, and proactive disclosure builds credibility while preserving control over the narrative. The disclosure should include four elements. First, a description of the AI governance framework as established in Part 4 of the program, including the seven principles and the eight-step methodology. Second, a description of the workflows currently in production (Finance Operations Copilot, Board Reporting) with their risk classifications and review patterns. Third, an example audit trail demonstrating the controls in operation. Fourth, the planned deployment roadmap for the remaining three workflows and the external review scheduled for the twelve-month anniversary of the governance charter. The deeper strategic point is that AI in finance is becoming a diligence topic in its own right at the Series C stage and beyond. A company that has done the work and can defend it is differentiated favorably from peer companies that have either avoided AI entirely or deployed it without governance. The diligence question becomes an opportunity rather than a risk.