AI for Controllers

Architecting Logic, Not Just Formulas

Safe, structured AI workflows for Finance in a BYOAI world

EFFORTLESS_METRICS // AI_FOR_CONTROLLERS

YOUR WORLD & AI TODAY

Your world

- Close: recs and variance memos
- Policy questions you research repeatedly
- SOPs you "mean to write" but never do
- Meetings where decisions get lost

AI today (ad-hoc)

- Excel formula help
- "Why is this not balancing?"
- "Googling" on ChatGPT

Useful, but not systematic

THE BOX WE'RE IN

- BYOAI is real: personal ChatGPT / Claude / Copilot are already in use
- PII still critical: no client/staff data, no ledgers, no raw exports
- No big AI platform launch: small workflow upgrades in tools you already use
- Everything today stays in that box — boringly compliant by design

WHERE THE DATA LIVES (AND WHERE IT DOESN'T)

Working workbook (real data)

- ERP / GL exports pasted into tabs
- Tabs that clean / map / calculate
- Close packs & audit files read from those

→ Contains real numbers + client/staff detail

Model copy (sample data only)

- Same tabs, headers, formulas as working workbook
- Inputs on dedicated / marked tabs
- Extract tabs → sample / rounded / dummy rows

→ Behaves like the working workbook — sample data only

HOW EXCEL SITS BETWEEN ERP AND OUTPUTS

Extract & transform (real data, no AI)

- **Extract:** ERP exports → input tabs
- **Transform:** formula-only tabs referencing inputs
- **Load:** results → templates / reports

→ Excel is your T in ETL; real PII stays on extract / input tabs

Model step (AI on sample copy)

- Copy workbook; replace extracts with sample rows
- AI sees headers, rules, formulas — not real data
- Ask for formulas, checks, structure, commentary

→ Copy stays PII-free

TWO SAFETY HABITS

Sanitise text

How you describe the situation

- **Describe > dump** — 3–5 scrubbed bullets, not an export
- Names → labels: “Client A (large CPG)”, “Senior Designer”
- No exports, contracts, HR/payroll records in prompts

> **If you wouldn't email it, don't paste.**

Sanitise workbooks

How you package the Excel

- Inputs live on dedicated tabs / clearly marked tabs
- Save As → "Model copy"; swap inputs for sample rows
- Only model copies with no client/ staff data ever go near AI

> **Upload a sample copy, never the working workbook.**

DEFENSE IN DEPTH

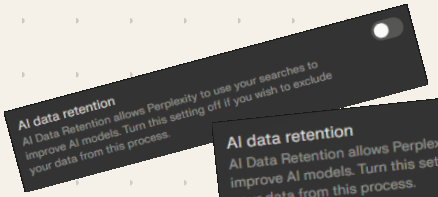

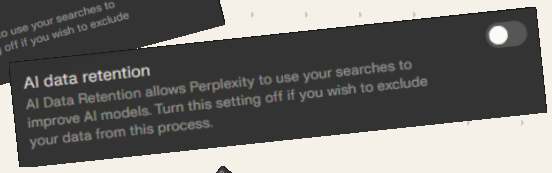

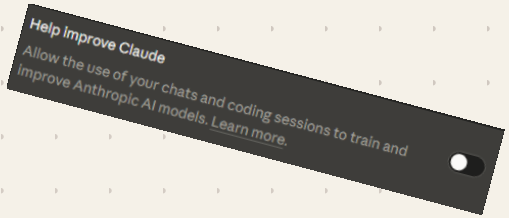

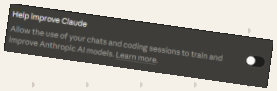

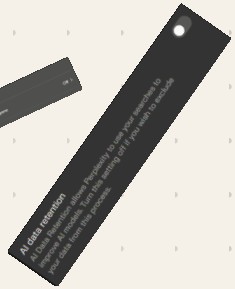

- Personal tools: turn off training / data-sharing in settings
- Not SOC-2, but reduces risk in case PII slips through
- Enterprise tiers (ChatGPT Team, Copilot, etc.) are SOC-2 — same bar as banking

Even with good habits, settings add a second layer.

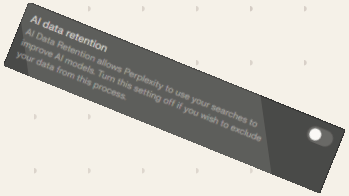

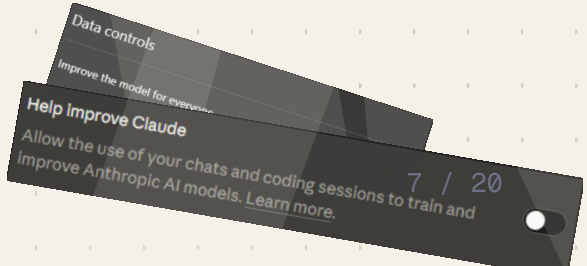

SETUP → WORK → CRITIC

Setup

What you give it

- Role: "You are a technical accountant for [company type]"
- Guidance: standards + Big-4 commentary on your topic
- Scenario: 3—5 scrubbed bullets; sample inputs only

> **Give it context before asking for output.**

Work & Critic

What you ask for

- Work: drafts, outlines, formulas — let it churn
- Critic: "Argue against this as a sceptical CFO / auditor"
- Skim Work; read Critic carefully

> **Context in, drafts out, then flip to hostile reviewer.**

LET THE MODEL CHURN, YOU READ ONCE

- Models output **12× faster** than you read — you're the bottleneck, not them
- You don't expect a junior's first draft to be good — same here
- Let it churn; one strong critic pass at the end

> **Use model speed to explore options, not to skip review.**

POLICY / STANDARDS QUESTIONS

Mental model

What the AI knows

- Junior technical accountant: read the standards, not your files
- Knows IFRS / ASPE / CRA + Big-4 commentary
- You own the treatment and sign-off

It's read the standards; it hasn't read your files.

When this helps

Recognise the moment

- "Does this meet criteria?" moments
- Ambiguous revenue / lease / tax scenarios
- Times you'd dig through standards or ping a colleague

Use it when you'd normally dig through PDFs.

POLICY / STANDARDS — HOW TO RUN IT

Setup

- "You are a technical accountant."
- "Summarise criteria & common traps for [topic] from IFRS / ASPE / CRA and Big-4 guidance."
- Run on Enterprise ChatGPT / Copilot, not personal BYOAI

Work then **Critic**

- Scrubbed 3—5 bullet scenario
- Questions to answer + evidence needed
- Plausible treatments and trade-offs
- "If we pick treatment A, argue against it as a cautious auditor / CFO."
- Use the arguments to harden your memo

PROCESS DOCS — STREAM → STRUCTURE

When to use it

Recognise the moment

- You're explaining a process out loud
- You "mean to" write the SOP but never do
- You have a messy email / Teams thread describing how it works

You've explained this before — capture it.

Inputs

What you feed it

- Phone recorder (on-device, training off)
- Teams / Zoom transcripts
- Rough bullets or a ranty email

Messy in, structured out — you fix the details.

PROCESS DOCS — FROM RANT TO SOP

Give it → Ask for

- Paste how the process actually runs
- An ordered outline
- A checklist (Step / Owner / Risk / Done)
- A 3—5 sentence summary for docs / onboarding
- Ask it to highlight contradictions or missing steps

You then fix

- Fix system names and paths
- Set thresholds and tolerances
- Add "watch-outs" and failure modes

MEETING NOTES THAT DON'T SUCK

What "good" looks like

The bar to hit

- Reliable: not "whoever remembered something"
- Findable / searchable; auditable (who, what, when)
- Note-owner is in the meeting, not in their notebook

Reliable, findable, auditable.

Why this matters

The payoff

- Most teams struggle; good notes keep focus after
- 1—2 note-owners per recurring meeting
- Transcript + AI capture → note-owner judges and publishes

AI captures; a human judges and publishes.

TRANSCRIPT → TEMPLATE → NOTE-OWNER

1. Transcript

- Teams / Zoom transcript (speaker IDs if available)
- Or phone recorder + non-training transcript

2. Template + AI

- Decisions, actions, owners, dates
- Open questions and risks
- Key terms / project context

Feed to Enterprise AI (SOC-2)

3. Note-owner

- Was actively listening during the meeting
- Uses draft to correct errors and fill gaps
- Sends notes to the right people
- Updates trackers / boards

RED TEAM — SECOND SET OF EYES

The problem

Why you need this

- Models tend to flatter whichever stance you give them
- Your reasoning has home-field advantage
- Important emails / memos are where gaps hurt most

> **Models flatter; your reasoning has home-field advantage.**

The goal

What you're after

- Treat the model as a hostile reviewer
- Use it to find holes before someone else does
- "What would a cautious CFO / auditor push back on?"

> **Find the holes before someone else does.**

RED TEAM — HOW TO RUN IT

Start → Sanitise

- Start from any draft: email, memo, commentary (AI-written or not)
- Remove names / IDs
- Keep structure and key facts / movements

In a fresh thread

- **A:** "Make this clearer and shorter without changing the facts."
- **B:** "Take opposite stance as a cautious CFO / auditor. Top 3 arguments + questions."
- **C:** "Rewrite to address those points without changing the facts." (optional)

Run in a fresh thread

ECONOMICS (THE 30-MINUTE RULE)

The math

What it costs vs. what you cost

- Enterprise AI seat: **~$40 CAD / month**
- Loaded controller time: ~$60—100/hour
- Breakeven: ≈ 30 min / month on work you'd do anyway

The honest question

Ask yourself

"Can this save ~30 min/month?"

- **If yes:** Seat justified on that workflow
- **If no:** Don't pretend it pays for itself

> **~30 min/month saved = seat pays for itself.**

NEXT STEPS — TRY ONE WORKFLOW ONCE

Pick one workflow that fits your world

- Policy / standards questions
- Stream → Structure for a process you own
- Meeting notes for a recurring series
- Red Team for one important email / memo
- Run it once on live work
- Did it save 10—30 min or catch something important?
  - If yes: keep it and standardise lightly
  - If no: cross it off and move on

Keep what works!

RECAP & Q&A

Key takeaways

- ETL mindset: PII in Extract/Load, sample copies to AI
- Safety habits: sanitise text + workbooks, toggle off training
- Setup → Work → Critic — let it churn, read Critic once
- Four workflows: policy, process docs, meeting notes, Red Team
- 30-minute rule: if one workflow saves that, seat pays for itself

> **Try something small and track if it saves time**