OPERATOR FIELD GUIDE NO. 02

FREE — NO SIGN-UP WALL

The Monday Money Hour
One prompt, one spreadsheet. Transactions in, a cash view and a written commentary out.

You connect your data once. From then on, one prompt categorises every transaction, files it into a daily tracker that rolls into a monthly view that rolls into a yearly P&L-style cash statement, compares it against your budget, and writes the commentary. In that order, in one run.

45 minutes	2 hours a week	No code	Free
TIME TO BUILD	TIME IT GIVES BACK	SKILL LEVEL	COST

DOWNLOAD THE TEMPLATE

cash-flow-tracker.xlsx — nine tabs, roughly 12,800 formulas, already wired. You never build a formula.

THE SHAPE OF IT

The spreadsheet does the arithmetic, the AI does the judgment

You connect your data once. From then on, one prompt categorises every transaction, files it into a daily tracker that rolls into a monthly view that rolls into a yearly P&L-style cash statement, compares it against your budget, and writes the commentary. In that order, in one run.

THE FLOW

```
CONNECTORS ─┐
            ├─▶ ONE PROMPT ─▶ Transactions ─▶ Daily ─▶ Monthly ─▶ Yearly
exports ────┘        │                                   │
                     │                    Budget ──▶ Variance
                     │                                   │
                     └────────── writes the commentary ◀─┘
```

THE DESIGN DECISION

Categorising is judgment. Summing is not. Most people get this backwards, hand the AI a pile of maths, and then wonder why the totals drift between runs.

Every rollup in the template is a live formula. Add a transaction in August and the August column, the yearly total, the variance and the closing cash all move on their own. The AI never touches a total.

STEP 1

Connect your data, so this stops being a chore

The version of this that survives is the one where you don't download anything. Set the connectors up once and the Monday job becomes "run the skill."

AIM FOR THIS CHAIN

LAYER	WHAT TO CONNECT	WHAT IT GIVES YOU
Bank and card	Your bank's scheduled statement export, or your accounting tool's bank feed	Transactions arriving without you fetching them
The landing zone	Google Drive or a shared folder, connected to your AI	The AI reads new files itself
The accounting tool	Xero, QuickBooks, Pleo, Ramp, Brex connector if one exists for your AI	Live data, no files at all
The workbook	Google Sheets connector, or Drive	The AI writes results straight into the tracker
The delivery	Slack or Teams connector	The commentary posts itself

Being realistic about banks. Very few banks offer a direct connector to an AI assistant today. The reliable pattern is one hop:

1. Set your bank and card provider to email or deposit a statement export on a schedule. Almost all of them do this. Weekly or monthly.
2. Point that at a single folder, `cash-flow/inbox`, in Drive or wherever your AI can read.
3. Connect that folder.

Now new data appears without you touching it, which is the actual goal. If your accounting tool already aggregates every account, connect that instead and skip the folder entirely.

> NAMING MATTERS MORE THAN YOU'D THINK
>
> Use `YYYY-MM-DD_account-name.csv`. The prompt reads the date range off the filename as a cross-check against the contents, which catches the single most common failure: running last week's file twice.

Connect the workbook too. If the tracker lives in Google Sheets rather than as a local file, the AI can write the transactions in directly and you never open a download folder again. This is the difference between a workflow you run for a month and one you run for a year.
No connectors available? Everything below still works. You upload two files instead. You lose the automation, not the output.

STEP 2

Set the template up, once

Open `cash-flow-tracker.xlsx`. Nine tabs.

TAB	WHAT IT IS
Setup	Base currency, reporting year, current month. Four cells.
Categories	Your buckets and how each maps into the P&L. **Edit this first.**
Transactions	The daily tracker. Every transaction, one row. The only tab data goes into.
Daily	Every date in the year, totalled by category, with a running cash balance.
Monthly	The same categories by month, arranged as a P&L.
Budget	Optional. Your budget by category and month.
Variance	This month and year to date, actual against budget.
Yearly	The full year against budget and against last year.
Commentary	Where each commentary gets logged, so you keep a history.

Three things to do before you run anything:

1. **Setup tab.** Base currency, reporting year, current month number, company name. Yellow cells only.
2. **Categories tab.** Fifteen categories are pre-filled with a P&L group each. Rename them to match how you actually think about spend. Keep the P&L group values exactly as written — the rollups match on that text.
3. **Monthly tab.** Type your 1 January cash balance into the January opening cash cell. Every later month carries forward on its own.

Budget is optional. Leave it empty and every variance column stays blank rather than breaking.

The one rule the whole thing depends on

> MONEY IN IS POSITIVE. MONEY OUT IS NEGATIVE.
>
> Every rollup assumes it. If your bank export uses separate debit and credit columns, or shows outflows as positive, the sign has to be fixed before it lands. The prompt below does that for you, but it's worth knowing why it matters: with this convention, gross profit is simply revenue plus cost of revenue, and net movement is the sum of everything. No sign gymnastics anywhere.

Two columns nobody should ever type into

On the Transactions tab, **column G (Amount in base)** and **column I (P&L group)** are formulas. G converts using your FX rate; I looks up the group from the Categories tab.

DELETING A ROW WIPES THEM SILENTLY

Delete a whole row and those formulas go with it, and that row silently stops counting. It'll look fine. It just won't be in any total.

So: write into columns A to F, H, and J to N. Never G or I. To clear a row, clear those columns individually rather than selecting the row and hitting delete. The prompt below is instructed accordingly, but if you're pasting by hand, this is the thing that will get you.

STEP 3

Pick your model

As of August 2026:

WHAT YOU'RE DOING	ON CLAUDE	ON CHATGPT
Building the skill	Fable 5	GPT-5.6 Sol
The weekly run	Sonnet 5	GPT-5.6 Terra
The commentary	Opus 5 or Sonnet 5	GPT-5.6 Sol or Terra
Auditing the setup	Fable 5, high effort	GPT-5.6 Sol
Never	Haiku	Luna

On the temptation to go cheap. Categorising transactions looks like the perfect job for the fastest, cheapest model. High volume, mechanical, repetitive. Don't. A miscategorised transaction that flows into a board number costs more than every token you'd save, and the cheap tiers degrade sharply past a few hundred rows — Luna sits around 41% on long-context recall against Sol's 91%.

THE DURABLE RULE, FOR WHEN THESE NAMES CHANGE

Build with the strongest model you have, run with the balanced mid-tier, never let the cheap tier near money.

STEP 4

The prompt

This is the whole workflow. One prompt, four stages, in order. It stops between stages if something looks wrong rather than pushing a bad number all the way through to a commentary you'd send to your CEO.

THE WEEKLY CASH FLOW PROMPT

```
You are running my weekly cash flow process. Work through all four stages in
order. If a stage fails a check, STOP and tell me. Never continue past a
failed check.

=== SETUP ===
WORKBOOK: [link to the Google Sheet, or the attached cash-flow-tracker.xlsx]
NEW DATA: [the connected folder, the accounting connector, or attached files]
BASE CURRENCY: [e.g. GBP]
PERIOD: [e.g. 4 to 10 August 2026]
MY ROLE: [e.g. Chief of Staff]
WHO READS THE COMMENTARY: [e.g. three founders, none of them in finance]

===================================================================
STAGE 1 — INGEST AND CHECK
===================================================================
Read every new transaction file, or pull from the connector.

Before anything else, report:
- each source, its date range, and its row count
- the date range you were expecting
- whether any transaction falls outside the stated period
- whether any transaction already exists in the Transactions tab
  (same date, same amount, same counterparty)

STOP AND ASK if: a date range does not match, a file looks like a repeat of
one already processed, or more than 5% of rows look like duplicates.
Never silently skip or silently include a duplicate.

Normalise every row to: date (ISO), description, counterparty, amount,
currency, FX rate to base.

SIGN CONVENTION, absolute: money in is POSITIVE, money out is NEGATIVE. If the
source uses separate debit and credit columns, or shows outflows as positive,
convert it. Then state explicitly which convention the source used and confirm
you converted it. Get this wrong and every number downstream is wrong.

===================================================================
STAGE 2 — CATEGORISE
===================================================================
Read the Categories tab. Use ONLY those category names, spelled exactly.

For each transaction assign a category and a confidence (high / medium / low).

RULES:
1. Below 90% confidence, use "Other or uncategorised". Never guess into a real
   category. A wrong category that looks right is worse than a blank one.
2. Never alter a date, amount or description. Copy them exactly.
3. Flag possible duplicates: same counterparty, same amount, within 3 days.
   Flag them, never delete them.
4. Flag anything more than 3x the typical size for its category.
5. Flag any counterparty that appears nowhere else in the workbook history.
6. Apply my saved rules below, they override your judgment:
   [paste your accumulated rules here, e.g. "AWS EMEA is Infrastructure and
   hosting", "anything from Fiverr is Payroll and contractors"]

===================================================================
STAGE 3 — WRITE TO THE WORKBOOK
===================================================================
Append every transaction to the Transactions tab, one row each, after the last
populated row.

WRITE ONLY into columns A to F, H, and J to N.
```

```
NEVER write into column G or column I. Those are formulas. Overwriting them
breaks the rollups silently.

Column mapping:
A Date · B Description · C Counterparty · D Amount · E Currency ·
F FX to base · H Category · J Confidence · K Flag · L Notes ·
M Source (the filename or connector) · N Entity

Do NOT touch the Daily, Monthly, Yearly or Variance tabs. They are entirely
formula driven and update themselves. Writing into them destroys the model.

Then confirm:
- how many rows you added
- the sum of column D for those rows
- that the sum matches the total of the source files, to the penny

If it does not match to the penny, STOP. Do not proceed to the commentary.

===================================================================
STAGE 4 — THE COMMENTARY
===================================================================
Only after Stage 3 reconciles.

Read the Variance tab, which is already calculated. Do NOT recompute anything.
Read the numbers as they are. If Budget is empty, skip every budget comparison
and write a movements-only commentary.

State first: the exact period of the actuals and the exact period of the
budget you are reading. If they do not match, stop and say so.

Then write the update. I will paste it straight into a leadership chat.

WRITING RULES:
- Short bullets. No paragraphs, no preamble, no sign-off.
- Number first, then the reason. Not the reverse.
- Every claim carries a figure. "Hosting up 22% to 14.2k", never "hosting
  increased".
- If you do not know why something moved, write "cause unknown, checking".
  NEVER invent a reason. This rule outranks all the others.
- No performance adjectives. No "strong", "concerning", "healthy". Give the
  number and let them judge.
- Maximum 8 bullets.

STRUCTURE:
**Cash** closing balance, and the movement on the week
**Against budget** the 2 or 3 biggest variances only, with value and percentage
**Worth a look** anything unusual, new, or newly large
**Clean** one line confirming what came in as expected, so they know it was
  checked and was fine

Then, separately, marked FOR ME NOT FOR THEM:
- every transaction that landed in "Other or uncategorised", so I can fix it
- everything you flagged
- the questions I should chase before sending this

Finally, append the commentary to the Commentary tab: date generated, period,
type, the commentary text, and the questions to chase.
```

Why it's built in stages

Each stage has a gate. Stage 1 stops on a date mismatch, which is how you catch running the same file twice. Stage 3 stops unless the total reconciles to the penny, which is how you catch dropped rows. Stage 4 refuses to compute anything, which is how the commentary can't contradict the spreadsheet.

> THAT LAST ONE IS THE IMPORTANT DESIGN DECISION
>
> The commentary reads the Variance tab rather than doing its own arithmetic. So the number in your Slack message is definitionally the number in the sheet. Let an AI compute the totals as well as narrate them and eventually those two things disagree, in public.

STEP 5

Your rules are the real asset

The first run will get a handful of calls wrong. Correct them out loud:

CORRECTING IT

```
Three corrections. "AWS EMEA" is Infrastructure and hosting, not Software.
"Stripe fee" lines are Banking and FX, not Professional services. Anything
from Fiverr is Payroll and contractors. Add these as permanent rules.
```

Then make it write them down:

EXTRACTING YOUR RULESET

```
Write out my complete categorisation ruleset as a numbered list, including
everything I have corrected. I want to save this and paste it into Stage 2.
```

After about four weeks, "Other or uncategorised" shrinks to almost nothing and stays there. **That accumulated ruleset, not the prompt, is what you actually built.** Keep it as its own document, because you will need to re-teach it — see the warning below.

STEP 6

Install it

On Claude

AS A PROJECT (START HERE)

1. Create a Project, `Weekly Cash Flow`.
2. Paste the Stage 1 to 4 prompt into the custom instructions.
3. Upload your ruleset and your voice file to the Project knowledge.
4. Connect Drive or Sheets, and Slack.
5. Monday: *"run the weekly cash flow."*

AS A SKILL

Folder `weekly-cash-flow`, file `SKILL.md`:

SKILL.MD FRONTMATTER

```
---
name: weekly-cash-flow
description: Run the weekly cash flow process end to end — ingest new bank and
  card transactions, reconcile the totals, categorise into the user's own
  buckets, write them into the Transactions tab of the cash flow tracker, then
  read the Variance tab and draft a budget versus actuals commentary in the
  user's voice. Use when the user asks to run their weekly cash flow, uploads
  bank statements or transaction exports, or asks for their finance update.
---
```

The full prompt underneath. Zip, rename to `.skill`, upload.

FULLY SCHEDULED

SCHEDULED TASK

```
Set this up as a scheduled task, every Monday at 08:00.
Check cash-flow/inbox for files newer than the last run. If there are none, DM
me "No new files, did not run" and stop. If there are, run all four stages.
Put the commentary in my inbox as a draft, and the FOR ME section in a DM.
Never post the commentary anywhere automatically.
```

> NEVER AUTO-POST A FINANCE COMMENTARY
>
> The draft step is your last look before a number reaches your leadership team.

On ChatGPT

Create a Project, `Weekly Cash Flow`, with the prompt in Project instructions and your ruleset as a Project file. Connect Drive so it reads the inbox folder itself. Use Tasks for the Monday nudge.

If your workbook is a local xlsx rather than a Sheet, ChatGPT can read the upload and hand you back the rows to paste. More clicking, same output.

STEP 7

Keep the step that makes you think

You can automate all of this. Bea deliberately doesn't:

> *"I still copy-paste the pivot table manually into our cash flow management tool... because I actually want to check it. And I like doing it. It could be done automatically, but I don't think everything is automatable should be automatable. I still like to look at those numbers."*
>
> BEATRICE ALIPRANDI — EP 11

The template gives you a natural place to keep that: **read the Variance tab yourself before you send the commentary.** Thirty seconds. It's where you notice hosting is up 30% and nobody mentioned it.

Automate the typing. Keep the noticing.

STEP 8

Make it fail loudly

Every gate above already stops rather than guessing. Add one more:

FAILURE INSTRUCTION

```
If the inbox has no new files, do not produce a report. DM me: "No new files
found, cash flow did not run this week." Send this even when nothing happened.
```

Silence and success look identical. A workflow that quietly skips a week is worse than none, because you'll believe the number you're looking at is current.

THE KNOWN FAILURE MODE, FROM THE EPISODE

The one time this went wrong for Bea, it *"was picking up the budget from the week before."* Not a crash. A quietly stale comparison, which is the dangerous kind. Stage 4 opens by stating both date ranges specifically to catch that.

Audit it

Monthly, on your strongest model:

SECURITY AUDIT

```
Review my weekly cash flow automation. What data can it read, what can it write
and where, whose credentials does it use, and what happens if a connector
breaks or a permission changes. Flag anything where access is wider than the
job needs. Then check the workbook itself: has anything overwritten a formula
in the Daily, Monthly, Yearly or Variance tabs.
```

That second question is the one that matters here. This is not theoretical — on the episode, Bea's own sweep found a real authentication flaw in a scheduled task, left over from reconnecting broken connectors. Reconnecting things can leave you worse off than before, and you cannot see it in the output.

THE GENERAL CADENCE

Every three prompts, check for bugs. Every ten, check for security.

FROM EXPERIENCE

Two warnings from people who've run this a while

Getting the voice right takes weeks, not an afternoon. Bea: *"trial and error. I'm actually very feedback oriented with Claude."* Ceci, who maintains a full written voice document, still puts hers at *"barely sixty percent to where I want it to be."* The shortcut is a voice file: five real updates you've sent, plus a banned-words list. That gets you further in an hour than a month of corrections.

Tuning can be undone. Bea moved her skills between Claude surfaces and the commentary quality collapsed:

> *"The transition was pretty rough... really, really rough. I was very sad for a long time... I had to train it for two or three more weeks."*
>
> BEATRICE ALIPRANDI — EP 11

A platform or model change can wipe weeks of tuning. Which is the real reason to keep your ruleset and voice file as separate documents: they're your backup.

BEFORE YOU CONNECT ANYTHING

The data question

This workflow involves your company's actual money.

- **Know where it goes.** Check your company's policy and whether your plan trains on your data — business and enterprise tiers generally do not, consumer free tiers may. One conversation, once, beats finding out later.
- **Strip what isn't needed.** The categoriser needs date, description, amount, currency. It does not need full account numbers, card numbers or counterparty addresses. Delete those columns. Thirty seconds, materially less exposure.
- **Consider aggregating payroll.** Individual salary lines are sensitive in a way the rest of this isn't. One line called Payroll works fine for a cash view.

HONEST LIMITS

When not to do this

- **Anything going to an investor, a board or a regulator** is a first pass, not a source of truth. A human reconciles before it leaves the building.
- **This is a cash view arranged like a P&L.** No accruals, no prepayments, no depreciation. It tells you what moved and when. Don't call it a P&L to your accountant.
- **Don't automate the review.** Categorising, flagging, totalling, first draft: automate freely. Deciding what the numbers mean is the reason you have the job.

TROUBLESHOOTING

When it misbehaves

"The totals don't match the source."

Stage 3 should have caught it. Ask: *"Show me the sum of column D for the rows you added against the source file totals, row by row, and find the discrepancy."*

"A row isn't in any total."

Column G or I got wiped. Copy the formula down from the row above, or the row silently doesn't exist.

"Everything is the wrong sign."

Stage 1's convention check was skipped. Ask what convention the source used and have it redo Stage 1 only.

"It's guessing categories."

Your bucket descriptions are too thin, or the 90% rule isn't being enforced. Add examples on the Categories tab and re-state the rule at the top.

"The commentary contradicts the spreadsheet."

It computed instead of reading. Reply: *"You are not permitted to calculate. Read the Variance tab as-is and quote it."*

"It invented a reason for a variance."

Reply: *"You stated a cause you cannot know. Replace it with 'cause unknown, checking'."* Then move that rule to the very top.

"Variance columns are all blank."

Budget tab is empty. That's the designed behaviour, not a fault.

THE POINT OF THIS

Automate in the right order

Five hours became three. The two hours that vanished were the typing. The three that remain are the thinking.

Most people automate in the wrong order. They build the beautiful dashboard and keep categorising by hand. Do it the other way round: let the machine sort and sum, and spend Monday morning looking at what it found.

> This guide came out of Top of the Ops, Season 1 Bonus — "Ride or Die AI Workflows", where Bea and Ceci went through every automation they actually run. The rest of the Operator Field Guides live on the Free Guides page.
>
> *Built something better? Tell us. We'll publish the improved version and credit you.*