

## FREE RESOURCE

# The QC Starter Kit

---

Make your quality control visible — from your very next review.

## WHY THIS EXISTS

When a human did the books, quality control was invisible because it was baked into the doing. Once AI produces the books, the data still arrives — but nobody vouched for it. Checking has to become a deliberate job, and the value isn't the quality control you claim; it's the quality control you can show. This kit is the first artifact that makes yours visible.

## WHAT'S INSIDE

### Correction Log

Excel template — 11 columns. One row per catch. Two worked examples included so you can see the level of detail that makes an entry useful.

### Exception Register

Excel template — 8 columns. When the same catch repeats for a client, promote it to a standing rule. This is the layer that compounds.

### This Guide

A capture prompt you paste into your AI tool, a worked example of feeding a correction back into context, and a 90-day cadence to make it stick.

## SECTION 1

# How to run it

Four steps — what each looks like in practice.

### Step 1. Capture during review, not after.

The single biggest failure mode is "I'll log it later" — you won't. The Correction Log entry happens while you're looking at the screen, in the same minute as the catch. The friction is real (it's another window). The fix is on page 3: a capture prompt your AI tool runs in 30 seconds, and it writes the row for you.

### Step 2. Promote repeats to the Exception Register.

A repeat is the same client plus the same pattern, twice. A one-off is noise; a repeat is a missing rule. The log is where individual catches live; the register is where the firm's knowledge actually accumulates. When you promote, give the rule a name a non-author can read in the field — "operating-line draws are not revenue," not "big deposit thing."

### Step 3. Close the loop.

The whole moat lives here. A captured rule that never makes it back into the AI's context is a rule you'll re-catch next month — and the month after that. Page 4 walks the full cycle on one real example, from the catch to the encoded context block. If you skip this step, the kit is just a diary.

### Step 4. Monthly 10-minute review.

Sort the Exception Register by Times Seen, largest first. Ask three questions: Which exceptions are still "Captured" but not yet "Encoded"? Which client shows the most corrections (they need the most context engineering)? What was your highest-severity catch this month (that's the story you tell when a client asks how you check the AI)?

## SECTION 2

# The capture prompt

Paste this into Claude or ChatGPT. After a catch, describe it in plain English. The prompt writes the Correction Log row for you.

You're helping me capture a quality-control catch on AI-produced bookkeeping. I'll describe the catch in plain English. Convert it to a single Correction Log row in this exact column order:

Date | Client | Engagement | AI tool / agent | What the AI produced | What was wrong (be specific) | Client-specific rule or context it missed | Correction made | Severity (High/Medium/Low) | Reviewer | Encoded back into AI context? (Yes/No/Add to AI context)

Be specific in the "rule it missed" column — name the actual client rule, not a generic category. If I haven't told you the date, severity, or reviewer, infer from context or ask one question. Output the row only, tab-separated, ready to paste.

Here's the catch: [paste your description]

### Variant for a team setting

Paste the prompt into a shared firm Claude project (or your firm's central AI workspace) so every reviewer gets identical row formatting. Standardized rows merge cleanly when you consolidate the team's logs for the monthly review — and they make the register's "Times seen" counter actually meaningful across reviewers.

## SECTION 3

# Closing the loop

One real catch on Northwind Café, walked all the way through to an encoded rule the AI won't repeat next month.

### A. THE CATCH

A \$40,000 deposit landed in Northwind Café's chequing account on May 26. The AI bookkeeper tagged it as Sales Revenue. It was actually a draw on the equipment loan we set up in April — same dollar pattern as a sale, completely different account.

### B. THE LOG ROW

Captured via the prompt on page 3 — one row in the Correction Log:

<b>Date</b>	2026-05-26
<b>Client</b>	Northwind Café
<b>AI tool</b>	AI bookkeeper
<b>What the AI produced</b>	Categorized \$40k deposit as Sales Revenue
<b>What was wrong</b>	It was an equipment-loan draw, not a sale
<b>Rule it missed</b>	Operating-line draws hit chequing as round-number deposits
<b>Correction made</b>	Reclassified to Equipment Loan (liability)
<b>Severity</b>	High — would have overstated revenue on the dashboard
<b>Encoded?</b>	Yes

### C. THE PROMOTION

This pattern had been caught twice before on Northwind. That makes it a rule, not a one-off. Promoted to the Exception Register as: "Operating-line draws are not revenue. Trigger: any deposit  $\geq$  \$10k with no matching invoice. Times seen: 3."

### D. THE ENCODING

The same rule, rewritten as a context block the AI tool will actually use:

Northwind Café — banking rules: Deposits from the operating line of credit land in chequing as round-number deposits (typically \$10k+). These are loan draws, NOT revenue. Code to Equipment Loan (liability) unless an invoice is matched. Flag any unmatched deposit  $\geq$  \$10k for review.

**Where the block goes — pick the one that matches your setup:**

- Per-client instructions field in your AI bookkeeping tool, OR
- A Claude Project for this client (paste in the project instructions), OR
- A shared firm SOP referenced by the reviewer's standing prompt.

Next month, the same kind of deposit lands. The AI now tags it as Equipment Loan. The catch goes from a 5-minute reclassification to a 30-second confirmation. That's the moat working.

## SECTION 4

# Your first 90 days

How to ramp from zero to a working QC habit without over-engineering it.

### WEEK 1

#### Start the log.

One reviewer, one client, every catch. Don't over-engineer it — no template tweaks, no team rollout, no monthly review cadence yet. The point is the habit, not the volume.

→ Expected output: 5–15 rows in the Correction Log. Zero entries in the Exception Register — you don't know what repeats yet.

### WEEKS 2–4

#### Watch for repeats.

Same reviewer, same client. The first Exception Register entries appear — promote anything you've caught twice. Encode the highest-frequency rules first (Section 3, page 4).

→ Expected output: 15–40 log rows. 2–5 register entries. 1–2 encoded rules. You'll notice the AI catches its own version of the encoded rule next time.

### MONTH 2–3

#### Expand and share.

Bring in a second reviewer or a second client. Standardize the capture prompt across the team (the variant on page 3). The monthly 10-minute review (Step 4) becomes a real meeting.

→ Expected output: a Correction Log that's measurably shorter month-over-month on the client you started with — because rules are being encoded faster than new ones are being caught. That's the leading indicator the moat is working.

The log is yours from day one — no one's permission needed. The compounding starts the day you make the first entry.