



CASE STUDY · FP&A WORKFLOW DESIGN

One sentence of context flipped the margin conclusion

A finance team ran the same contribution-margin analysis twice — same data, same model, same day. The only variable was what the analyst told it. The two runs recommended opposite actions.

SUBJECT

H1 2026 contribution margin review

WORKFLOW

10 steps · 2 human checkpoints

DATASET

324 rows · 12 accounts × 5 products × 6 months

BUILT & RUN

August 2026

About the data. This case is based on a real client engagement, recreated end-to-end with **synthetic data** for illustration. "Meridian Components," its accounts, products and every figure in this document are fabricated. The dataset was generated from a fixed seed so the numbers genuinely reconcile — the variance bridge ties to a zero residual — but no real company data appears here. The workflow, the design decisions and the sequence of events are real.

THE SITUATION

Half-year close. Contribution margin came in at 38.89% against a 40.00% plan — a miss of 111 basis points. Contribution dollars, meanwhile, were **above** plan by \$40,460 on 11,427 extra units.

That combination is the kind of thing that turns a review into an argument. The rate says one thing, the dollars say the opposite, and the person who has to explain it to leadership has a week.

CONTRIBUTION MARGIN	CONTRIBUTION \$	REVENUE	VOLUME
38.89%	\$15.49M	\$39.83M	260,311
Plan 40.00%	Plan \$15.45M	Plan \$38.62M	Plan 248,884 units
-111 bps	+\$40,460	+\$1.21M	+11,427

Synthetic dataset at account × product × month grain.

The constraint that defines the problem

The actuals export carries eight fields: period, account, channel, region, product, units, price per unit, variable cost per unit. There is **no discount field**. No rebate field. No commercial-agreement flag.

Which means that in this data, a deliberate price concession, a competitive loss and a mix shift are **mathematically indistinguishable**. No amount of analysis resolves that. No better model resolves it either. It is a missing column, and the only place the missing information exists is in the memory of whoever was in the meeting.

That constraint is what the workflow was designed around.

THE WORKFLOW


Ten steps, built as discrete skills so that each one has a single job and a declared **force** — whether its output is locked, interpretive, or human.

01 load-actuals DETERMINISTIC	02 load-budget DETERMINISTIC	03 compute-margin DETERMINISTIC	04 decompose-variance DETERMINISTIC	05 read-the-variance PROBABILISTIC
06 Frame the question HUMAN	07 diagnose-margin-gap PROBABILISTIC	08 Checkpoint HUMAN	09 update-context DETERMINISTIC	10 build-margin-report DETERMINISTIC

■ Deterministic — the math cannot vary

■ Probabilistic — the two steps that interpret

■ Human in the loop



The workflow as built. Each node carries its own trigger phrases, declared inputs, and instructions; the two *HITL* nodes — **Frame the question** and **Checkpoint** — are the only two steps that stop and wait for a person.

Facts, then orientation, then frame

Steps 1–4 are locked. They read the files, reject bad rows with a reason for each, compute contribution at every cut, and split the variance into volume, mix, price and cost. Step 4 refuses to hand anything forward unless the bridge reconciles exactly — **budget + volume + mix + price + cost = actual**, residual zero.

CONTRIBUTION BRIDGE VS PLAN	USD	BPS OF REVENUE
Budget contribution	15,449,497	
Volume effect	+709,332	
Mix effect	-7,788	-2
Price effect	-543,843	-137
Cost effect	-117,241	-29
Actual contribution	15,489,957	
Residual	0	0

Price carries the miss. What the bridge cannot say — and does not try to — is why price moved.

The step that refuses to conclude

Step 5 is the design decision the whole workflow turns on. Its job is to **widen the analyst's view, never narrow it**. It describes what stands out, lists several situations consistent with the same numbers — competitive pressure, a deliberate concession, mix within accounts, rebate timing, a change in buyer mix — presented **unranked**, names precisely which field is missing from the export, and then stops and asks.

It is explicitly forbidden from saying which reading is most likely, or from pointing at the two accounts where the effect is concentrated.

That restraint is counter-intuitive, so it is worth stating plainly why it is there. If step 5 hands the analyst a conclusion, the human step that follows collapses into a rubber stamp — an expensive way to say yes. The workflow would produce the same answer with or without the person in it, and the answer would only ever be as good as what the data already contained.

THE EXPERIMENT

The workflow was then run twice against the identical dataset, on the same day, with the same model. **One variable changed: what the analyst typed at step 6.**

RUN A — THE THIN FRAME

WHAT THE ANALYST SAID

“Just tell me what’s driving it.”

WHAT CAME BACK

Contribution margin is 111 bps below plan. Price is the dominant driver at -137 bps. Realized price declined across the portfolio.

What I cannot tell from the data: whether this reflects competitive pressure, deliberate commercial decisions, or customer mix. The data shows realized price fell; it does not show why.

Correct. Honest. Well-written. And it gets you nowhere in a business review. The natural reading — we are losing pricing power across the book — sends someone to tighten discount policy.

RUN B — THE REAL FRAME

WHAT THE ANALYST SAID

“In March we gave our two largest distributors a 4% off-invoice discount to hold volume against a competitor buying share. It was decided in a meeting. It is not in the ERP. It touches those two accounts only.”

WHAT CAME BACK

Isolating those accounts: they ran on plan through February, then from March realized price stepped down -4.23% — matching the 4% off-invoice — while their volume rose +15.4%. That single action is **95% of the company's entire price variance.**

The rest of the book held price flat. Its -36 bps is an even variable-cost drift of about \$0.40 a unit.

The portfolio-wide erosion was an artifact of blending. A deliberate -286 bps on two accounts, averaged with a -36 bps cost drift on the other ten, produced a -111 bps headline that looked company-wide and was not.

Same workflow, same data, same model, same day. Both runs are reproducible from the seeded dataset.

What the split actually looks like

ACCOUNT	CONTRIBUTION \$	VS PLAN	CM% ACTUAL	CM% PLAN	CM BPS	UNITS VS PLAN
Ironwood Industrial	3,753,478	+59,345	38.09%	40.15%	-206	+6,166
Calloway Holdings	3,059,628	+49,883	38.05%	40.08%	-203	+4,979
All other accounts (10)	8,676,852	-68,768	39.56%	39.92%	-36	+282
Total	15,489,957	+40,460	38.89%	40.00%	-111	+11,427

Two accounts breach the materiality threshold (75 bps / \$250k). Both move more than 200 bps. Everything else is quiet.

The two conclusions, side by side

Run A concludes the company is losing pricing power and should tighten discount governance. Run B concludes the company deliberately bought volume in two accounts and needs to evaluate whether the trade paid.

Those are opposite management actions. One of them restricts the commercial team; the other asks whether to extend what the commercial team did. Nothing in the data distinguished them. The distinguishing fact lived in a conversation, an email thread, and a handshake.

THE JUDGMENT

Step 8 does not ask “does this look right.” That question can be answered yes without thinking. It asks two questions that do different work: a decision the data genuinely cannot settle, and the rule that would let the workflow reach that decision unaided next time.

The decision here was genuinely close. On the concession accounts the margin **rate** fell 286 bps from March while contribution **dollars** rose \$118,285 against plan. Rate and dollars point in opposite directions — which is exactly where a decision sits rather than a calculation.

“Keep it — I'd make the same call again. We conceded about \$519K of price and the two accounts came back with \$109K more contribution than plan, so in dollars the trade paid. But the reason I'd keep it isn't the \$109K. It's what those accounts are. Ironwood and Calloway are our two largest distributors. A competitor was buying share, and if we'd held price to protect the rate, the risk wasn't a slower quarter — it was losing a distributor. You don't get those back.”

The judgment, recorded verbatim at the checkpoint

Then the extension — the part no model produces, because it had not happened in any dataset yet:

“That said, I'm not comfortable. I can't prove the discount bought the volume, and neither can you. What I know is that I was on the QBR with Ironwood last week and they're consolidating suppliers next year — so this volume isn't as safe as these numbers make it look.”

Context supplied at the checkpoint, held in no system

Note what the checkpoint captured that a sign-off would not have: a decision, the reasoning behind it in the decider's own words, a business fact that exists in no system, a scope and an expiry date, and an executable rule for next time.

WHAT THE CYCLE PRODUCED

Step 9 wrote all of it back to a context file that steps 5 and 7 read before they run again. That is the return arrow — what makes this a loop rather than a line, and what makes the next cycle start higher than this one did.

A decision with an explicit expiry

Keep the concession. It is **fixed-term** — it runs to year-end and is renegotiated at renewal. Recorded with the instruction not to model it into H2 as a permanent rate, which is precisely the error the next forecast would otherwise have made.

Two rules the workflow now applies unaided

First: when an account's realized price drops more than 200 bps in a month while that account's volume rises, treat it as a commercial action rather than erosion — flag it, name the account, and ask. Second: stop reporting blended CM% as the headline when one account moves more than 200 bps against the rest of the book. Show the split.

A data request with an owner

FP&A to Commercial: add a discount / commercial-agreement flag to the actuals export, with start and end dates. Until that field exists, the two rules above are the interim workaround. This is the durable fix — the analysis was never the bottleneck, the missing column was.

An open item carried forward

The even -29 bps variable-cost drift across the rest of the book is unrelated to the concession, and a prior-quarter rule said the input-cost exposure would be revisited at the half-year. It was not this cycle's framed question and was explicitly not decided — so it is recorded as open rather than quietly closed.

A management one-pager

A single self-contained page — headline tiles, the contribution waterfall, the account split, and three commentary columns. It leads with the business event rather than the variance, and contains no reference to the process that produced it. Reproduced in full overleaf.

WHAT THIS CASE SHOWS

The framing step is where the value enters

Both runs were competent. The difference between them was not analytical quality — it was that one of them knew about a meeting. Placing the human **before** the diagnosis rather than only after it is what let that knowledge into the analysis at the point where it could still change the question being asked.

A checkpoint that collects agreement is not a checkpoint

The reason step 5 is forbidden from naming the likely cause is that the alternative degrades quietly. A system that presents a conclusion and asks for approval will get approval — and will produce identical output whether or not anyone was paying attention. The absence of a recommendation is what forces the human contribution to be real.

The blind spot moved up a level

Before the frame, the open question was *why did price fall*. After it, the open question is *whether that volume survives the concession lapsing* — a harder and more useful question, and one the checkpoint surfaced by admitting that the causal link is unprovable from the data. The cycle did not only produce a better answer. It produced a better next question, and wrote down enough that the next cycle does not start from zero.

Most of the pipeline should not be able to vary

Six of the ten steps are deterministic — loading, validating, computing, decomposing, writing memory, and rendering the report. Only two steps interpret, and only two are human. Every figure in the final deliverable traces back to the single step that owns the arithmetic; no downstream step is permitted to calculate anything. That is what makes the interpretive steps safe to trust: they can be wrong about meaning, but they cannot be wrong about the numbers.

Synthetic data notice. This case study is based on a real client engagement and has been recreated with entirely synthetic data for illustration. Meridian Components does not exist; the accounts, products, figures and quoted remarks are fabricated for the reconstruction. The dataset was generated from a fixed seed so the arithmetic genuinely reconciles and the workflow can be re-run and reproduced. The workflow design, the sequence of events, and the conclusions drawn about the method reflect the real engagement.

Management report, as produced by step 10

Reproduced as the workflow rendered it. Every figure traces to the deterministic compute and decompose steps; nothing about the process appears on the page — this is what leadership receives.

Contribution Margin Review

Meridian Components · H1 2026 (January–June)


Prepared by FP&A

Date 7 Aug 2026

Currency: USD

CONTRIBUTION MARGIN % 38.89% Plan 40.00% –111 bps · by design	CONTRIBUTION \$ \$15.49M Plan \$15.45M +\$40,460 vs plan	REVENUE \$39.83M Plan \$38.62M +\$1.21M vs plan	VOLUME (UNITS) 260,311 Plan 248,884 +11,427 vs plan
--	---	--	--

CONTRIBUTION BRIDGE VS PLAN — WHERE THE \$40K NET CAME FROM



ACCOUNT SPLIT — ACCOUNTS PAST THE MATERIALITY THRESHOLD (75 BPS / \$250K)

ACCOUNT	CONTRIBUTION \$	VS PLAN \$	CM% ACTUAL	CM% PLAN	CM BPS	UNITS VS PLAN
Ironwood Industrial	3,753,478	+59,345	38.09%	40.15%	-206	+6,166
Calloway Holdings	3,059,628	+49,883	38.05%	40.08%	-203	+4,979
All other accounts (10)	8,676,852	-68,768	39.56%	39.92%	-36	+282
Total	15,489,957	+40,460	38.89%	40.00%	-111	+11,427

Only Ironwood and Calloway breach materiality (both > 200 bps). The other ten accounts are shown as one summary row: a steady -36 bps driven by an even variable-cost drift, flat on price.

WHAT HAPPENED

In March we gave our two largest distributors, **Ironwood Industrial** and **Calloway Holdings**, a **4% off-invoice discount** to hold volume against a competitor who was buying share — a commercial decision made in a meeting, not visible in the ERP export.

The data confirms it exactly: those accounts ran on plan through February, then from March realized price stepped down **-4.23%** while their volume rose **+15.4%**. That single action is **95%** of the company's price variance.

The rest of the book held price flat; its -36 bps is a small, even variable-cost drift (~\$0.40/unit) across every account.

WHAT IT MEANS

Rate and dollars point opposite ways, and that is the whole story. The concession cost **-286 bps** on those accounts from March but added **+\$118k** of contribution — volume on ~40% margin more than paid for the price given up.

The company-wide **-111 bps** headline is a blend, not a problem: a deliberate -286 bps on two accounts averaged with a -36 bps cost drift on the other ten.

Decision: keep the concession. The rationale is strategic, not the \$109k: Ironwood and Calloway are our two largest distributors, and holding price to protect the rate risked losing a distributor to a competitor buying share — a loss that doesn't come back.

WHAT TO WATCH

We cannot prove from the data that the discount *bought* the volume — the discount and the volume simply arrive together. So the +15.4% is less safe than it looks.

- **Fixed-term:** runs to year-end, renegotiated at renewal — do **not** model into H2 as a permanent rate.
- **Action — FP&A → Commercial:** add a discount / commercial-agreement flag to the export, with start & end dates. Today a concession, a competitive loss, and a mix shift are indistinguishable.
- **Carry forward:** the even -29 bps cost drift across the book — frame the resin / input-cost exposure explicitly next cycle.