



Commercial Diagnosis — Initial Audit

Meta Ads Channel

Prepared for: a UK personalised gifting & jewellery retailer (name withheld)

Scope: Meta Ads, full account, 90-day performance window (1 Jun – 24 Aug 2026)

Prepared by: GOAS — James Moody

THE DIAGNOSIS

Three things you need to know

01

Your 200% ROAS target is hiding the difference between growth and demand capture.

02

Your two live ad sets look 6.6x apart on cost per purchase — but they're not measuring the same thing.

03

£10,787 of historic spend went through a 15%-off code that produced a worse cost per purchase than the rest of the account.


THE BIGGEST GAP

No agreed New Customer CAC ceiling exists. Until it does, there's no reliable definition of what "profitable acquisition" actually means for this account.

WHY THE NUMBERS DISAGREE

The number is lying to you — here's how

£20,008.83 spent. 4,932 purchases. Judged against a single 200% ROAS target, that looks reasonably healthy. It isn't the full picture.



Two of your live ad sets are running on genuinely different measurement models — not different quality, different arithmetic.

Scale — New Sales ad set

Measured on Incremental attribution: only counts purchases that wouldn't have happened without the ad. Deliberately strict. Will always look more expensive per purchase than standard attribution, even with identical creative and identical audience quality — because it's not inflating its own number.

Capture — Standard Perso

Measured on standard 7-day click / 1-day view attribution: credits any purchase in that window, including people who'd have bought anyway. Looks efficient. Some of that efficiency is real; some of it is the measurement being generous to itself.

Judged against one blended target, real growth can look like the underperformer and lose budget, while spend that's mostly claiming credit for existing demand looks like the star and gets more. That's not hypothetical — it's live in this account today.

THE FINDINGS

Five findings

Problem, evidence, why it matters, what we'd do — quantified everywhere the data supports it, and flagged honestly everywhere it doesn't.

FINDING 01

A single ROAS target is judging campaigns it can't fairly compare

● **LEARN** *not quantifiable — the risk is invisibility itself*

THE PROBLEM

The account targets one blended 200% ROAS across the whole Meta account, regardless of what job each campaign is doing or how it's measured.

THE EVIDENCE

New Sales ad set (Incremental attribution) and Standard Perso (standard attribution) are both judged against the same 200% line, despite measuring fundamentally different things.

WHY IT MATTERS

Real net-new growth can be starved of budget for looking "inefficient", while spend that's simply better at claiming credit for existing demand gets rewarded with more.

WHAT WE'D DO

Set a separate return threshold per campaign job, matched to its attribution method. Revisit it regularly — this drifts as campaigns run and audiences overlap.

COMMERCIAL IMPACT: UNQUANTIFIED

FINDING 02

The two live ad sets are being compared on the wrong basis

● **LEARN** *live now, not historic*

THE PROBLEM

New Sales ad set looked 6.6x less efficient than Standard Perso on cost per purchase — a gap that reads as a targeting or creative failure.

THE EVIDENCE

£11.58 vs £1.76 cost per purchase. The gap is very largely explained by attribution methodology, not audience quality: Incremental attribution is deliberately stricter.

WHY IT MATTERS

Acting on this gap at face value — e.g. cutting New Sales ad set's budget — would mean punishing the campaign for being honestly measured, and rewarding the one that isn't.

WHAT WE'D DO

Stop comparing these two on blended CPA. Judge New Sales ad set against a contribution-margin ceiling instead. Judge Standard Perso against whether it still has a live purpose.

COMMERCIAL IMPACT: UNQUANTIFIED

FINDING 03

New Sales ad set has no target CAC to be judged against

● **SCALE** *directional only, not verified*

THE PROBLEM

There's no agreed ceiling for what a new customer is allowed to cost — so £11.58 per purchase can't be called good or bad with real confidence.

THE EVIDENCE

A directional check (Meta's ROAS-implied revenue, VAT removed, your stated 78% margin) puts contribution margin at roughly £15.90/order against an £11.58 cost per purchase — a rough £4.32 cushion. Not verified against real order data; doesn't include non-COGS costs; implied £20.39 AOV looks low for this retailer's typical basket.

WHY IT MATTERS

Without this ceiling, every future decision about this campaign's budget is a guess dressed up as an analysis.

WHAT WE'D DO

Build a verified NCAC ceiling from real order economics. Until then, treat this campaign as directionally healthy, not confirmed profitable.

COMMERCIAL IMPACT: DIRECTIONAL — ~£4.32/order cushion

FINDING 04

223 ads sit dormant across 35 zero-spend ad sets

● DEFEND score: 6/10

THE PROBLEM

The account is carrying a large amount of structure from the previous agency that isn't doing any job at all.

THE EVIDENCE

35 ad sets, £0 spend across the full 90-day window, 223 ads sitting inside them.

WHY IT MATTERS

Dormant structure left live creates confusion for anyone working in the account, and risk of accidental reactivation with stale creative or offers attached.

WHAT WE'D DO

Archive the zero-spend ad sets before building anything new on top of the account.

COMMERCIAL IMPACT: LOW EFFORT, LOW RISK, WORTH DOING FIRST

FINDING 05

Historic reliance on a 15%-off code was not even the most efficient spend

● **RECOVER** score: 3/10

THE PROBLEM

The previous agency built the account heavily around a 15%-off code (ANGEL15) — margin given away as standard practice, not as a targeted exception.

THE EVIDENCE

£10,787 — 54% of all recorded spend — ran through ANGEL15-tagged ads, at a £4.50 cost per purchase. The remaining 46%, without the discount, achieved a lower £3.63 cost per purchase.

WHY IT MATTERS

The discount wasn't just a margin risk — it was measurably the less efficient choice. There's no performance case for rebuilding around it.

WHAT WE'D DO

Don't rebuild the new campaign structure around a blanket discount code. The data argues directly against it.

COMMERCIAL IMPACT: QUANTIFIED — £10,787

COMMERCIAL OPPORTUNITY

What this audit actually found

£10,787

historic spend identified as inefficient relative to the rest of the account — the one number in this report we're confident enough in to lead with.

QUANTIFIED

£10,787 of historic spend at a worse cost per purchase than the rest of the account (Finding 5).

DIRECTIONAL

~£4.32/order contribution cushion on New Sales ad set — real signal, not yet a verified number (Finding 3).

UNQUANTIFIED

The measurement framework itself is misjudging which campaigns deserve budget — the size of that cost isn't visible yet, and that invisibility is the risk (Findings 1 & 2).

We don't manufacture a score where the data doesn't support one. Two of five findings here don't have a clean £ figure — that's reported honestly, not smoothed over.

WHAT TO DO

Priorities

NOW

Archive the 35 dormant ad sets. Check the creative on the four expired-promo ads before that campaign is ever reactivated.

NEXT

Give Standard Perso a current, deliberate goal, or retire it. Set attribution-matched return thresholds instead of one blended ROAS target.

INVESTIGATE

Build a verified New Customer CAC ceiling from real order economics. This is the highest-value open question in the account, and the one everything else should ultimately be judged against.

WHAT'S NEXT

The questions worth answering properly

This Initial Audit has identified the highest-priority questions in this account. Answering the biggest one properly — a verified New Customer CAC ceiling, built from real order economics rather than a directional estimate — is deeper work than a £499 review is built to carry.

That's the natural next step: not a bigger version of this report, but the piece of work this report has shown is actually needed.