

Meta Advertising Audit & 90-Day Growth Plan

Sample deliverable — client identifiers redacted

Client: [REDACTED] — DTC personalised keepsake jewellery brand

Account reviewed: [REDACTED]

Media spend: approximately \$6,500 per month

Situation: working with a retained agency

Review window: 90 days

Prepared by: Serafinity

1,964 people added a product to their cart in the last 90 days. 390 of them reached checkout. 192 bought.

The advertising is working. The checkout is not. This document quantifies the difference and sets out how to capture it.

About this sample. This is a real audit delivered to a real client, published with their written permission and anonymised at their request. Brand name, domain, account and dataset identifiers, campaign and asset names, and product names have been redacted. Currency figures are rounded. Every percentage, ratio, benchmark and conversion rate is exactly as delivered. Nothing in the analysis has been softened, removed or improved after the fact.

Contents

1. Scope and method — what was examined, and what could not be verified
2. Primary finding: only 19.9% of carts reach checkout — the number that decides this account
3. Diagnostic summary — seven findings ordered by revenue impact
4. Account health and delivery baseline — Opportunity Score 91/100
5. Spend distribution across campaigns
6. Funnel conversion rates, stage by stage — every step, benchmarked
7. What is happening at checkout — the personalisation hypothesis and seven tests
8. Retargeting: currently inactive
9. Creative inventory and performance — six findings including an under-funded winner
10. Permanent discount structure
11. Tracking and signal quality — duplicate purchase events and a stalled CAPI
12. Competitor activity in the US market — 422 active ads against 8
13. Unclaimed positioning — six positions available now
14. 90-day remediation plan — week by week
15. Budget scaling thresholds — gated
16. Creative concept territories — eight territories, each tied to a finding
17. Creative production model
18. Seasonal planning
19. Projected outcomes, timelines and risks — including what looks worse first
20. Reporting and working cadence
21. Summary of findings — the whole audit on one page

1. Scope and method

This audit is based on primary data pulled directly from the client's Meta advertising account and the associated pixel dataset. No figure in this document is estimated unless explicitly labelled as such.

What was examined

- All 28 campaigns in ad account [REDACTED], across the 90 days ending in the review window.
- Campaign, ad set and ad-level delivery and conversion metrics, including spend, impressions, reach, frequency, CTR, CPM, CPC, ROAS, cost per purchase, and every funnel event from landing page view through purchase.
- Video engagement data at ad level: plays, 25% retention, 75% retention, and ThruPlay, for every ad with meaningful spend.
- Meta's Opportunity Score and its full recommendation set for the account.
- Pixel dataset [REDACTED], including browser and server event freshness, event coverage, and hour-by-hour event volumes for the trailing seven days.
- The live Meta Ad Library for the United States, filtered to active ads only, across multiple category and outcome keyword sets.
- The [REDACTED] storefront: product range, price architecture, promotional structure, and product imagery.

What could not be verified

Three inputs sit outside the advertising account and were not available at the time of writing. Where a conclusion depends on one of them, that dependency is stated in the text rather than buried.

- Contribution margin after cost of goods, shipping and payment processing. Every SKU on the site carries a permanent discount, so the margin position is not inferable from list prices.
- Backend revenue from Shopify for the same 90-day window, which is the only way to test Meta's attribution against reality.
- Repeat purchase rate and customer lifetime value. For a keepsake brand with a natural second-purchase cycle, this materially changes what a profitable first-order ROAS looks like.

These are the first three items we would request on day one of an engagement. They sharpen the targets; they do not change the diagnosis.

2. Primary finding: only 19.9% of carts reach checkout

Across the 90-day window, 1,964 add-to-cart events were recorded against Meta traffic. Only 390 of those progressed to an initiated checkout. That is a progression rate of **19.9%**.

Healthy ecommerce accounts convert 50% to 65% of add-to-carts into initiated checkouts. This account converts one in five.

Add to cart to purchase: 9.8%. Category norm is 25% to 35%. The account is converting carts into orders at roughly one third of the rate the category expects.

Why this is the finding rather than one of several

Everything upstream of the cart is performing acceptably. CPMs sit at \$12.97 on the primary campaign, which is competitive for the parenting and gifting audience. Frequency is 3.13 over 90 days, nowhere near saturation. Click-through rate is 1.36%. Meta's own Opportunity Score for the account is 91 out of 100, which means the delivery mechanics, placements, budget structure and bidding are broadly aligned with platform best practice.

In other words, the media buying is not the problem. The account is successfully buying attention, successfully converting that attention into product interest, and successfully persuading people to add jewellery to a cart. Then four out of five of those people leave.

What the leak is worth

SCENARIO	ATC TO CHECKOUT	PURCHASES	REVENUE	ROAS
Current (90 days)	19.9%	192	\$42,650	2.17
Checkout at 30%	30.0%	290	\$64,380	3.27
Checkout at 40%	40.0%	386	\$85,720	4.36
Checkout at 55% (norm)	55.0%	531	\$117,910	6.00

All four rows assume identical ad spend, identical creative, identical targeting and the existing 49.2% checkout-to-purchase rate. The only variable changed is the single step between adding to cart and beginning checkout.

The 3x ROAS goal is reached at a 30% cart-to-checkout rate. That is still twenty points below the category norm. It does not require better ads, a bigger budget, or a new agency doing clever things in Ads Manager. It requires fixing one step in the purchase path.

This is also the explanation for a pattern the business has almost certainly noticed already: increasing spend has never produced a proportional increase in revenue. More spend pushes more people into a cart that leaks at the same rate. The leak scales with the budget.

3. Diagnostic summary

#	FINDING	IMPACT	SECTION
1	Cart-to-checkout progression at 19.9% against a 50-65% norm	Critical	2, 6, 7
2	Purchase event double-firing; CAPI stale for four days	Critical	11
3	Retargeting switched off entirely against a 36,000-customer base	High	8
4	Ad sets optimising for purchase count, not purchase value	High	4, 14
5	Creative concentration: 3 assets carry 48% of spend	Medium	9
6	Permanent sitewide discount removes all promotional headroom	Medium	10
7	Engagement campaigns bought 2,095 clicks and zero purchases	Contained	5

The one-paragraph version

This is a competent Meta account with a broken measurement layer sitting on top of a leaking checkout. Meta reports 2.17 ROAS across the window, but the pixel is firing purchase events twice, which means the true figure is materially lower and must be confirmed against Shopify before any target is set. Underneath the measurement problem, the funnel loses four out of every five people between the cart and the checkout. Fix the tracking so decisions are made on real numbers, fix the checkout so the traffic already being bought converts, restart retargeting against a customer base that is currently being ignored, and the 3x goal is reachable inside 90 days without increasing spend.

4. Account health and delivery baseline

Meta's Opportunity Score for the account is 91 out of 100. This is an account-level score measuring alignment with platform best practice on structure, bidding, placements and audience configuration.

A 91 is high. For context, accounts with genuine media-buying problems typically score in the 50s and 60s and generate a long list of structural recommendations. This account generated seven, and six of them are the same recommendation repeated across different campaigns.

METRIC (90 DAYS)	VALUE
Total spend	\$19,660
Meta-attributed revenue	\$42,650
Meta-reported blended ROAS	2.17
Purchases	192
Cost per purchase	\$102.40
Average order value	\$222
Implied monthly spend	\$6,550
Opportunity Score	91 / 100

Note on budget

Spend across the window averages \$6,550 per month, against a stated budget of approximately \$8,000. Either delivery has been throttled, spend has been paused for part of the window, or a second account carries the balance. This is worth confirming, because the scaling plan in section 15 is built off actual delivered spend rather than nominal budget.

What Meta is actually recommending

Six of the seven recommendations say the same thing: duplicate the ad sets to use the Maximise Value of Conversions performance goal with the Purchase event, for an estimated 7% ROAS improvement. This applies to nine separate campaigns, including the full 17-ad-set test cluster.

The account is optimising for purchase count rather than purchase value. On a price ladder that runs from \$120 to \$550, this is not a rounding error. Meta's delivery system is being told that a \$120 necklace and a \$440 bracelet are the same outcome, so it will reliably chase the cheaper one. The 7% lift Meta estimates is conservative for a catalogue with this much price spread.

The seventh recommendation suggests enabling Shops ads on the primary prospecting campaign for an estimated 6% reduction in cost per conversion. Given the checkout findings in section 7, routing some traffic through Meta's native checkout rather than the site's own configurator is more interesting than the 6% headline suggests. It is included as a test in section 14.

5. Spend distribution across campaigns

The account contains 28 campaigns. Eight delivered spend in the window. Two of those carry 90% of it.

CAMPAIGN	SPEND	ROAS	PURCH.	CPP	CTR
Prospecting (CBO)	\$11,640	2.49	124	\$93.90	1.36%
Test cluster (ABO)	\$6,110	1.87	60	\$101.81	1.39%
Retargeting (paused)	\$830	0.91	4	\$206.47	0.98%
Cold broad, copy (paused)	\$365	1.25	3	\$121.68	1.33%
Engagement, copy (paused)	\$320	—	0	—	5.63%
Scale (paused)	\$190	—	0	—	1.35%
Sales (paused)	\$130	7.66	1	\$130.22	1.01%
Cold broad (paused)	\$75	—	0	—	5.67%
Total	\$19,660	2.17	192	\$102.41	

Observations

- **Prospecting is the account.** It carries 59% of spend, 65% of purchases, and returns 2.49 ROAS against a 2.17 blended average. It is the strongest structure in the account and the correct foundation to build on.

- **The test cluster** spent \$6,110 to return 1.87, dragging the blended figure down by roughly 0.2x. It contains 17 ad sets, which for that spend across 90 days means each ad set received an average of \$4 per day — well below the threshold at which Meta's delivery system can exit the learning phase and optimise meaningfully.
- **Twenty campaigns delivered no spend** in the window. Some are legacy tests dating back nearly two years. This is housekeeping rather than a performance issue, but a 28-campaign account with 8 live campaigns makes weekly reporting harder than it needs to be and increases the risk of accidental reactivation.
- **The Sales campaign** shows 7.66 ROAS on a single \$997.50 purchase from \$130 of spend. This is one order, not a trend. It is listed for completeness and should not be read as a signal.

Structural fragmentation

The test cluster is the clearest example of a pattern worth naming. Seventeen ad sets sharing \$6,110 over 90 days means no individual ad set ever accumulated enough conversion signal to leave learning. Under the current auction architecture, where Meta's delivery system needs volume through a single optimisation surface to model effectively, splitting a modest budget across seventeen ad sets is the most reliable way to guarantee that none of them work. Consolidation is addressed in section 14.

6. Funnel conversion rates, stage by stage

STAGE	VOLUME	STEP RATE	BENCHMARKVERDICT
Clicks	21,374	—	—
Landing page views	11,697	54.7%	70–85%Below
View content	8,161	69.8%	60–80%Normal
Add to cart	1,964	24.1%	10–20%Strong
Initiated checkout	390	19.9%	50–65% Critical
Purchase	192	49.2%	60–75%Below

Reading the table

Two steps are performing well. View-content sits comfortably in range, and add-to-cart at 24.1% of product views is genuinely strong, above the category norm. People want this product. The creative is doing its job and the product pages are doing theirs.

Three steps need attention, in descending order of severity.

Add to cart to initiated checkout: 19.9%. This is the critical failure and the subject of section 7. It costs more revenue than every other issue in this document combined.

Initiated checkout to purchase: 49.2%. Against a 60–75% norm, roughly a fifth of the people who begin checkout are being lost during it. Notably, the two main campaigns diverge sharply here: prospecting converts 44% of checkouts, while the test cluster converts 73%. A thirty-point spread between two campaigns hitting the same checkout is unusual and points at either audience quality differences or, more likely given the findings in section 11, event attribution noise.

Click to landing page view: 54.7%. Nearly half of paid clicks never register as a landing page view. Some of this is normal — people tap and bounce before the page loads, and iOS privacy measures suppress a portion of events. But 54.7% is at the low end. The two candidate causes are mobile page load speed and pixel fire timing. Both are testable in an afternoon and both are included in the week one work in section 14.

Clean-data confirmation. The cart leak is not an artefact of ad tracking. Site-wide pixel data for a four-day period when, as section 11 explains, the data was firing cleanly shows 122 add-to-carts against 20 initiated checkouts across all traffic sources. That is 16.4%, which is if anything worse than the paid-traffic figure. The leak is in the storefront, not in the advertising.

7. What is happening at checkout

This brand does not sell jewellery. It sells jewellery engraved with a specific child's date of birth, weight, height and time

of birth. That personalisation is the entire value proposition and the reason the brand can charge \$220 where the category charges \$40.

It is also, almost certainly, where the money is being lost.

The hypothesis

A customer sees an ad, likes the piece, and adds it to the cart. At or immediately after that point, the purchase path asks them to supply exact birth data: date, time, weight in pounds and ounces, length in inches. This is a fundamentally different cognitive task from buying a necklace.

- **It requires information most people do not have in working memory.** Time of birth and birth weight usually live on a document in a drawer, not in someone's head.
- **It carries permanence anxiety.** This is engraved metal. Getting a digit wrong is not returnable, and buyers know it. Hesitation at an irreversible step is rational.
- **It arrives at the worst possible moment.** Meta traffic is overwhelmingly mobile and overwhelmingly impulse-led. The gap between an emotional impulse and a data-entry task is where conversion dies.
- **It is frequently a gift purchase.** A grandparent, aunt or friend buying for a new parent may not have the birth details at all, and has no easy way to obtain them without ruining the surprise.

The data supports this. Add-to-cart is strong, which means desire is being created successfully. The collapse happens precisely at the transition into the checkout flow. That is the signature of a friction problem, not a desire problem or a price problem. If price were the barrier, the drop would appear earlier, at product view.

Test sequence

#	TEST	HYPOTHESIS	EFFORT
1	Move personalisation after payment	Capture the order first, collect engraving details by email or a post-purchase form. Removes the friction from the conversion path entirely.	Medium
2	Add "I'll add details later" option	Same principle, lighter build. Lets the impulse complete and defers the data task.	Low
3	Save-and-return cart with reminder	Lets someone leave to find the birth certificate and come back without losing the cart.	Low
4	Persistent proof preview	Show a live rendering of the engraving as details are entered. Directly targets permanence anxiety.	Medium
5	Gift mode	Recipient supplies their own details after the gift is purchased. Unlocks the entire gifting audience currently blocked at this step.	High
6	Mobile checkout audit	Field count, keyboard types, tap targets, autofill behaviour on the actual devices the traffic arrives on.	Low
7	Express checkout at cart	Shop Pay, Apple Pay and Google Pay surfaced before the personalisation step rather than after.	Low

A note on ownership. Tests 1 through 7 are storefront changes, not advertising changes. We will specify them, provide the exact copy, and measure the result. Implementation sits with whoever holds the Shopify theme. We are telling you this before you have signed anything because the alternative — taking the account, running the same traffic into the same checkout, and reporting a modest ROAS improvement in month three — would be the easier pitch and the wrong one.

8. Retargeting: currently inactive

The brand states 36,000 families served. The only retargeting campaign in the account is paused. While it ran during the window it spent \$830 and returned 0.91 ROAS. It lost money.

METRIC	RETARGETINGREAD
Spend	\$830
ROAS	0.91Below breakeven
Purchases	4
Cost per purchase	\$206.472.2x the account average
Frequency	6.70Saturation
CTR	0.98%Lowest in the account
Reach	8,473Very small pool

Why it failed, and why that is good news

A frequency of 6.70 against a reach of 8,473 people means the same small group saw the same ads nearly seven times in the window. CTR of 0.98%, the lowest in the account, is what audience exhaustion looks like. The campaign was not badly conceived; it was pointed at too few people with too little creative variety, and it burned out.

This matters because pausing it was the right immediate call and the wrong permanent one. Retargeting is normally the cheapest revenue in any ecommerce account. For this brand it should be cheaper than average, for three reasons specific to the product.

- The natural repurchase cycle is real and predictable. A second child means a second piece. Grandmothers, aunts and godparents buy the same design for the same baby.
- Mother's Day and Q4 give two annual moments when a past customer has an obvious reason to buy again, and when a brand that already holds their child's birth details has a structural advantage over every competitor.
- The 1,574 people who added to cart and never checked out in the last 90 days are, by definition, the highest-intent audience the brand has. They are currently receiving nothing.

The immediate opportunity. 1,574 people added a product to their cart in the last 90 days and did not buy. At the account's \$222 average order value, converting even 10% of that group represents roughly \$35,000 in recovered revenue — comparable to 80% of everything the account produced in the entire window.

The rebuild is specified in section 14. The short version: segment by behaviour rather than lumping all site visitors together, cap frequency, and give each segment a distinct creative angle rather than re-serving prospecting assets.

9. Creative inventory and performance

Forty ads carried meaningful spend in the window. The table below covers those above \$400.

ASSET	SPEND	ROAS	CPP	CTR	HOOK
Asset A — baby with ring, video	\$5,690	2.22	\$98.04	1.49%	34.6%
Asset B — product on background	\$2,360	2.58	\$84.20	1.43%	9.3%
Asset C — ring film, lullaby score	\$1,470	2.58	\$73.54	1.97%	20.2%
Asset D — product still	\$1,100	2.84	\$78.96	0.78%	4.0%
Asset E — testimonial, "mom like me"	\$1,080	2.42	\$98.52	1.11%	16.7%
Asset F — best sellers mashup	\$900	2.20	\$74.74	1.14%	4.9%
Asset G — product on background, copy	\$590	1.88	\$84.72	1.58%	9.3%
Asset H — heart pendant, video	\$480	7.46	\$47.71	1.55%	15.7%
Asset I — birthday angle	\$430	3.20	\$106.97	0.78%	3.7%

Hook rate = 25% video retention as a share of video plays. Category norm is 20–30%.

Finding 1: there is an under-funded winner

Asset H returned 7.46 ROAS at a \$47.71 cost per purchase, less than half the account average, on \$480 of spend. That is the single most efficient asset in the account by a wide margin, and it received 2.4% of the budget.

The caveat is honest: ten purchases is a small sample and the figure will regress. But an asset performing three times the account average deserves a deliberate scale test rather than being left at \$5 a day. This is a week one action.

Finding 2: concentration risk

Three ads carry 48% of all spend, and one ad carries 29% on its own. This is fragile in two directions. If that asset fatigues, roughly a third of account delivery degrades at once with nothing tested and ready behind it. And under an auction that rewards creative diversity within a single campaign structure, a narrow asset pool caps how much the delivery system can find distinct pockets of the audience.

Finding 3: the asset library is smaller than the ad count suggests

Forty ads is misleading. One asset appears under five different ad IDs across four campaigns; another appears four times. Stripping duplicates, the account is running roughly twelve to fifteen genuinely distinct creative concepts, and several are variations on the same shot.

Finding 4: hook rates are mostly below par

Asset A achieves a 34.6% hook rate, which is genuinely good and explains why it earns its budget. But four of the nine assets above sit below 10%, and two below 5%. Asset D retained 4.0% of 88,898 plays. That means 96% of people who started the video were gone before the three-second mark. These assets are being carried by whatever happens after the click, not by the creative itself.

Finding 5: discount-led creative underperforms emotional creative

Two ads in the account lead with the discount in the ad name and the creative. The first returned 0.76 ROAS. The second produced no purchases at all. Meanwhile the emotionally framed assets — the baby with the ring, the lullaby-scored ring films, the testimonial — carry the account. This is a consistent and useful signal, and it connects directly to section 10.

Finding 6: the product photography is AI-generated, and the filenames say so

Live product pages are serving image files whose names are unmodified AI export filenames, complete with the generator name and timestamp: [REDACTED] and [REDACTED].

For most categories this would be a minor housekeeping note. For this one it is a strategic problem. The brand's entire proposition is that the piece is real, handcrafted from recycled precious metals, and permanently marked with a specific child's actual measurements. AI-generated product imagery undercuts every one of those claims for any customer who looks closely, and it removes the strongest asset the brand has: photographs of the real thing on real hands, at real scale, with real engraving depth. It is also plausibly contributing to the checkout hesitation described in section 7. If the product images do not feel real, committing to a permanent engraving feels riskier.

The remedy is straightforward and is built into the creative plan in section 17.

10. Permanent discount structure

Every product on the site carries a struck-through list price and a sale price. Not most products. Every product.

PRODUCT	LIST	SALE	DISCOUNT
Necklaces (from)	\$150	\$120	20%
Rings	\$180-\$280	\$144-\$224	20%
Bracelets	\$225-\$365	\$180-\$292	20%
[REDACTED] bracelet	\$550	\$440	20%
[REDACTED] cuff	\$400	\$320	20%

Three consequences

First, the promotional lever is spent. When Mother's Day and Q4 arrive — the two moments in the year when this

brand should be able to drive urgency — there is nothing left to offer. A 20% discount that has been running continuously is not an offer; it is the price. Going deeper means cutting into margin that is already reduced, and the audience has been trained to expect the reduction as standard.

Second, urgency cannot be manufactured credibly. Countdown timers and limited-time framing on a permanently discounted catalogue read as untrue to any customer who has visited before, and this is a considered-purchase category where people visit more than once.

Third, the discount reduces contribution margin on every single order, which lowers the ROAS required to break even and narrows the room to bid competitively. Without margin data we cannot say whether 3x is comfortably profitable or barely breakeven. That distinction determines whether the scaling plan in section 15 is aggressive or cautious, and it is the first number we would ask for.

Recommendation

This is a pricing decision, not an advertising one, so we are flagging it rather than prescribing it. But the options are worth stating: restore true list pricing and reserve genuine discounting for two seasonal windows; or hold current pricing as the real price, remove the struck-through comparison, and compete on craft and meaning rather than on a discount nobody believes. The second is more consistent with a \$220 keepsake brand and with the creative territories in section 16.

Notably, the account's own data already argues for this. The two discount-led ads in the account returned 0.76 ROAS and zero purchases respectively, while the emotionally framed assets carry the entire account. This audience is not buying on price.

11. Tracking and signal quality

This section changes how every other number in this document should be read. It should be actioned before any budget decision is made.

Finding: the Conversions API stopped four days ago

The pixel dataset shows a browser event firing at 13:15 on the final day of the review window — current. The server-side event stream last fired at 21:44 four days earlier. The Conversions API has been silent since.

Finding: purchase events are firing twice

Hour-by-hour event data for the trailing seven days shows a clear and repeating pattern. Through the day the server stream stopped, the Purchase event consistently fired in pairs while its sibling events fired once.

hour	Purchase	NewCustomerPurchase	PurchaseProfitReading
Day 1, 20:00	2	1	1Duplicated
Day 2, 18:00	2	1	1Duplicated
Day 3, 14:00	2	1	1Duplicated
Day 4, 11:00	2	1	1Duplicated
Day 5, 08:00	1	—	—Clean
Day 6, 12:00	1	—	—Clean
Day 7, 06:00	1	—	—Clean
Day 8, 10:00	1	—	—Clean

Purchase fires twice; NewCustomerPurchase and PurchaseProfit, which describe the same order, fire once. Then on the day after the server stream stopped, Purchase drops to one and stays there.

The conclusion is difficult to avoid: the browser pixel and the Conversions API are both reporting purchases, and event deduplication is not working. Both streams are being counted.

A second corroborating signal

In several hours during the affected period, AddPaymentInfo exceeds InitiateCheckout — in one hour, four AddPaymentInfo events against two InitiateCheckout events. Payment information cannot logically be entered by more

people than began checkout. This is a funnel-order violation and is consistent with uneven duplication across event types.

What this means for the numbers in this document. The 2.17 ROAS, the 192 purchases and the \$42,650 revenue figure are what Meta reports. If purchase events have been duplicating for a meaningful part of the 90-day window, the true figures are lower, potentially substantially. Every performance number in this audit should be treated as an upper bound until reconciled against Shopify. The funnel diagnosis in sections 2 and 6 is unaffected, and is confirmed independently by the clean-period data in section 6.

A third signal, from volume

Site-wide pixel data for the clean four-day period shows approximately ten purchases across four days from all traffic sources — roughly 2.5 per day, or about 75 per month. Meta claims 192 purchases over 90 days, which is around 64 per month. Meta is therefore claiming credit for something close to 85% of every order the store receives, including organic, direct, email and returning customers. Some over-attribution is normal on a 7-day-click, 1-day-view window. This is beyond normal.

What else is configured

The dataset is otherwise reasonably set up. First-party cookie tracking is enabled. Eight events are firing: PageView, ViewContent, Search, AddToCart, InitiateCheckout, AddPaymentInfo, Purchase, and SubscribedButtonClick, plus the custom NewCustomerPurchase, ReturningCustomerPurchase and PurchaseProfit events. The custom customer-type events are a genuinely good piece of setup that is currently not being used for optimisation. They make new-customer acquisition cost measurable directly, which most accounts cannot do.

Remediation

- Restore the Conversions API connection and confirm server events are flowing.
- Implement event_id deduplication across browser and server streams. This is the actual defect.
- Reconcile 90 days of Meta-reported purchases against Shopify order counts and revenue. Establish the true ROAS.
- Re-baseline all targets against the reconciled figure before any budget change.
- Verify the Purchase event passes accurate order value rather than a static or rounded amount, given the \$120-\$550 spread.
- Once clean, enable value-based optimisation as Meta recommends in section 4, and put the NewCustomerPurchase event to work for acquisition-cost reporting.

12. Competitor activity in the US market

A live sweep of the Meta Ad Library, filtered to active ads in the United States, across several keyword sets.

SEARCH	ACTIVE US ADSREAD
personalized jewelry mom	422Crowded, commoditised
new mom keepsake baby birth gift	8Effectively unclaimed
personalized birth stats necklace new mom	0Nobody is there

That contrast is the single most useful competitive fact in this document. The generic personalised-jewellery auction has 422 active advertisers in it. The specific territory this brand actually occupies — birth statistics, new motherhood, the keepsake of a particular child — has almost nobody. Of the eight results in the second search, most were irrelevant. Only one was genuinely in the space.

Who is actually buying the auction

ADVERTISER TYPE	APPROACH	CREATIVE PATTERN
Volume leader	Category incumbent	Keyword-stuffed product titles; identical creative duplicated across 6–8 simultaneous ad IDs
Direct imitator	Follower	Same spec-sheet headline format; birthstone and birth-flower SKUs
Premium minimal	Upmarket	Short, restrained product descriptors
Price-led	Discount	"Personalized essentials from \$36" repeated across carousel cards
Sentiment-led	Gift-object framing	"To My Daughter I Am So Proud Of You"
Configurator-led	Build-your-own	"Design your meaningful family ring today"
Narrative DR	Story-first	"I Had One Day Left to Find Mom's Gift. Then I Found This."
Social proof	Scale claim	"More than 3M People Trust Us"

The dominant creative pattern, and why it is beatable

The overwhelming majority of active ads in this category use the product title as the ad headline. A representative example runs to roughly thirty words naming the material, the personalisation type, the product category, and every possible recipient from mother to grandmother.

That is search-engine copy pasted into a social feed. It is written for a keyword index, not for a person scrolling. It names every possible product attribute and every possible recipient in the hope that one lands. It creates no feeling whatsoever.

Two advertisers in the entire sweep write like direct-response marketers. One leads with a story — a real moment, a real deadline, a real relief. The other leads with proof at scale. Neither of them is a jewellery brand competing directly, and both of them are outperforming the category's copy conventions by a wide margin.

Structural observations

- The volume leader duplicates the same creative across six to eight concurrent ad IDs. This is a placement and delivery-share strategy, and it means the raw ad count in this category substantially overstates the number of distinct creative ideas in market.
- The visible price band in the category sits around \$30–\$80. This brand sits at \$120–\$550. Competing on the same keywords and the same spec-sheet copy against a \$36 alternative is a losing auction. Competing on a different emotional territory is not.
- Nobody in the sweep is using birth statistics as a creative device. Not the date, not the weight, not the time. The most emotionally powerful and most defensible asset this brand owns is entirely absent from the category's advertising.

13. Unclaimed positioning

Six positions available in this auction now, each supported by a finding above.

1. The birth statistics themselves

7 lb 4 oz. 3:42 in the morning. 20 inches. These numbers are the most emotionally loaded data a parent will ever encounter, and no competitor is using them. A creative built around the numbers rather than around the jewellery is unlike anything else in the feed and is impossible for a \$36 name-necklace factory to copy credibly.

2. Real craft against AI imagery

The category has moved wholesale to generated product imagery — and, per section 9, so has this brand's own storefront. That creates an opening for whoever moves first in the other direction. Hands, workbenches, engraving tools, metal being finished, the piece at true scale on a real wrist. In a category where everything looks synthetic, evidence of a human making the thing is differentiation that also happens to be true.

3. The gift-giver, not the mother

Every competitor addresses the wearer. Nobody addresses the grandmother, the husband, the sister or the best friend — the people who actually buy most keepsake jewellery, and who currently cannot complete a purchase because they do not have the birth details. This is simultaneously the biggest creative gap and, via the gift mode test in section 7, the

biggest checkout fix.

4. Permanence as the reason to pay more

The category sells personalisation as a feature. This brand can sell permanence as a value: this is recycled solid precious metal, it will outlast the childhood it records, and it will be handed on. That argument justifies \$220 against \$36 in a way no feature list can.

5. Proof at scale

36,000 families, and coverage in [REDACTED]. Only one advertiser in the sweep uses social proof as a primary hook, and it is not a jewellery brand. This is sitting unused.

6. The second piece

Nobody in the category advertises to their own past customers around the natural repurchase moment. A brand that already holds a child's birth details has an unmatched advantage in the second-baby, grandmother and Mother's Day conversations. Section 8 covers why this is currently worth nothing and how that changes.

14. 90-day remediation plan

Days 1-7: stop the bleeding, fix the instruments

Nothing else in this plan is trustworthy until the measurement layer is fixed. Week one is deliberately unglamorous.

- Restore the Conversions API connection and implement event_id deduplication.
- Reconcile 90 days of Meta purchases against Shopify. Establish true ROAS and re-baseline every target in this document.
- Run the mobile checkout audit and instrument the cart-to-checkout step so the drop-off is observable in real time.
- Scale Asset H out of its \$5/day corner into a properly funded test.
- Archive the twenty dormant campaigns.
- Diagnose the click-to-landing-page gap: page load speed and pixel fire timing.

Days 8-30: fix the checkout, restart retargeting

- Deploy checkout tests 2, 3, 6 and 7 from section 7 — the low-effort set — and measure cart-to-checkout weekly.
- Specify and scope test 1, deferred personalisation, which is the highest-value change in this document.
- Rebuild retargeting as three behavioural segments: cart abandoners (1-14 days), product viewers (1-30 days), past customers (30-365 days). Distinct creative per segment. Frequency capped at 3.
- Consolidate the 17-ad-set test cluster into two or three properly funded ad sets.
- Migrate prospecting ad sets to value optimisation per Meta's recommendation, once tracking is verified clean.
- Shoot the first real-photography creative batch. Six to eight assets.

Days 31-60: build the creative engine

- Launch creative territories 1 through 4 from section 16 as a structured test against the current control set.
- Replace the AI product imagery on the highest-traffic product pages with real photography.
- Deploy test 5, gift mode, if the earlier checkout tests confirm the friction hypothesis.
- Establish weekly creative cadence at three to five new assets per week.
- Begin the first scaling step from section 15, provided the gate conditions are met.
- Test Shops ads on the primary prospecting campaign per Meta's recommendation.

Days 61-90: scale on evidence

- Scale spend against the ladder in section 15, one step at a time, each gated on efficiency holding.
- Launch territories 5 through 8 and retire fatigued assets on a defined schedule.
- Build and load the Q4 creative slate — see section 18.
- Introduce broad campaign structure with high creative diversity as the primary scaling vehicle.

- Deliver a reconciled 90-day performance review against Shopify, not against Meta.

15. Budget scaling thresholds

Scaling from \$8,000 to \$50,000 a month is not a decision to be taken once. It is five decisions, each of which should be made by the data rather than by ambition.

The gates below are conditions, not a timetable. If a condition is not met, the step does not happen, and the previous step continues until it is.

STEP	MONTHLY SPEND	GATE CONDITION
0 — Baseline	\$8,000	Tracking reconciled with Shopify. True ROAS established. This gate must clear before any increase.
1	\$12,000	Cart-to-checkout above 30%. ROAS at or above 2.5 on reconciled data. Retargeting live and profitable.
2	\$18,000	ROAS holding at or above 2.8 for 14 consecutive days. At least 8 distinct creative concepts live. Frequency below 3.5.
3	\$28,000	ROAS at or above 3.0 sustained 21 days. Cart-to-checkout above 35%. New-customer CPA stable or falling.
4	\$40,000	ROAS at or above 3.0 at step 3 volume. Creative production sustaining 12+ new assets monthly. CPM increase under 20% versus baseline.
5	\$50,000	All step 4 conditions held for 30 days. Contribution margin confirmed positive at scale.

Rules that govern the ladder

- No step increases spend by more than 50% at once. Larger jumps reset the learning phase and reliably destroy efficiency for a week or more.
- A minimum of 14 days at each step before the next is considered. Ecommerce data is noisy below two weeks and the considered-purchase cycle here is longer than average.
- Any step that fails its gate steps back one level rather than holding position. Sitting at a spend level that is not working costs money every day.
- ROAS at every gate means reconciled ROAS measured against Shopify, not Meta-reported ROAS.

Why we are writing the gates down. Most agencies say they will scale when performance allows and then define "allows" after the fact. Written gates mean the decision to spend more of your money is made by evidence you can check, and it means we can be held to it.

16. Creative concept territories

Eight territories, each addressing a finding in this audit. Each is a territory rather than a single ad, meaning it supports multiple executions and can be tested and scaled independently.

Territory 1 — The Numbers

The birth statistics alone, held on screen. 7 lb 4 oz. 3:42 AM. 20 inches. No product for the first several seconds. Then the reveal: the numbers are engraved on the piece.

Why it works: section 12 established that no competitor uses birth data as a creative device. These numbers carry more emotional weight for a parent than any product photograph, and they cannot be copied by a factory selling name necklaces. This is the most differentiated concept available to the brand.

Territory 2 — Made By Hand

Workbench footage. Engraving tools. Metal being finished. The specific moment a number is struck into the surface.

Why it works: directly answers the AI-imagery problem in section 9 and the permanence anxiety in section 7. Evidence that a human makes the thing is the strongest available argument for \$220 over \$36, and it is currently absent from the

entire category.

Territory 3 — The Person Buying It For Her

Addressed to the grandmother, the husband, the sister. Not to the mother wearing it.

Why it works: every competitor speaks to the wearer. This audience is unaddressed in the category and, per section 7, is also structurally blocked at checkout. Pairing this creative with gift mode converts a creative gap and a checkout fix into the same win.

Territory 4 — Longer Than The Childhood

Permanence as the value proposition. Recycled solid gold and silver. The piece outlasts the baby it records, and gets handed on.

Why it works: reframes price as durability rather than as a premium, which is the only framing that survives comparison against a \$36 alternative. Also supports the pricing recommendation in section 10.

Territory 5 — 36,000 Families

Proof at scale. Real customer photography, real wrists, real engraving. Press logos used as reinforcement rather than as the hook.

Why it works: only one advertiser in the competitive sweep leads with social proof, and it is not a jewellery brand. Section 13 identified this as sitting unused.

Territory 6 — What You Forget

The emotional truth that specific details fade. You remember the day. You forget the weight, the exact time, the length. The piece remembers for you.

Why it works: creates the need before presenting the product, which is the classic direct-response sequence and the thing the category's spec-sheet copy entirely fails to do.

Territory 7 — The Second One

Built for retargeting past customers. Second baby, grandmother's version, matching pieces.

Why it works: section 8 established that the repurchase audience is receiving nothing. This territory exists to be pointed at the 36,000-customer base and at the Mother's Day window.

Territory 8 — Still Time

Deadline-driven urgency built on shipping cut-offs rather than on discounts.

Why it works: section 10 established that discount urgency is not available to this brand. Shipping deadlines are true, they are genuinely urgent, and the narrative advertiser's performance in the competitive sweep suggests deadline framing converts well in this category.

Testing structure

Territories 1, 2, 3 and 6 go first — they address the largest gaps and the highest-severity findings. Each launches with three executions in 9:16 and 1:1. Territories 4, 5, 7 and 8 follow in days 31–60. Every territory competes against the current control set rather than against a clean slate, so results are measured against what the account already achieves.

17. Creative production model

The scaling ladder in section 15 requires twelve or more new assets a month by step four. That volume is where most ecommerce accounts stall, and it is worth being explicit about how it gets produced, because "you will need more creative" is not a plan.

What we produce

- Concept development and scripting for all eight territories, including hook variants.
- Full creative briefs: shot lists, on-screen copy, timing, safe zones for each placement, and audio direction.
- Static and motion asset production, including generated imagery where it is appropriate — which, per section 9, means anything except the product itself.

- All ad copy: primary text, headlines, descriptions, and CTA variants for testing.
- Format adaptation across 9:16, 1:1 and 4:5.

What we need from the brand

Less than most agencies would ask for, and specifically not a photographer on retainer.

- **One real product photography session.** Pieces on real hands and wrists, in natural light, at true scale, with visible engraving depth. This is the single highest-value asset the brand can create and it replaces the AI imagery identified in section 9. One day of shooting supplies months of creative.
- Access to existing customer photography and any user-generated content already held.
- Approval on brand voice and any claims requiring verification.

Cadence

PHASE	NEW ASSETS / MONTHFOCUS
Days 1–30	6–8Territories 1–3, retargeting rebuild
Days 31–60	10–12Territories 4–6, real photography rollout
Days 61–90	12–16Territories 7–8, Q4 slate, winner iteration

Review is staged rather than batched: concepts are agreed before assets are produced, and assets are reviewed individually rather than as a monthly dump. This keeps revision cycles short and prevents a month of production going in the wrong direction.

18. Seasonal planning

This brand has two peak windows, not one. A 90-day plan starting in late summer runs directly into the first and must have the second built before it arrives.

Q4 — November and December

Gifting season, and the audience skews heavily toward the buyer rather than the wearer, which makes Territory 3 and the gift mode checkout fix commercially urgent rather than merely interesting. A grandmother who cannot complete a purchase because she does not have the baby's birth weight is a lost order in October and a lost order multiplied across the whole of December.

- Creative slate loaded and tested by 1 November. Testing new concepts during peak CPMs wastes money.
- Shipping deadline creative (Territory 8) built and scheduled in advance, with dates confirmed against fulfilment capability.
- Retargeting budget weighted upward through December, when intent is highest and the cart abandoner pool is largest.
- Expect CPMs to rise 30–60% versus the summer baseline. The scaling gates in section 15 should be read against seasonally adjusted expectations rather than paused reflexively.

Mother's Day — March through early May

The single largest moment of the year for this category, and the one where the brand's structural advantage is greatest. A business that already holds a customer's child's birth details can speak to that customer in a way no competitor can.

- Ramp begins mid-March, not in April. The category's CPMs escalate sharply from the first week of April.
- Territory 7 pointed at the full past-customer base — grandmother's version, matching pieces, the second child.
- Per section 10, urgency for this window has to come from shipping deadlines and genuine scarcity rather than from a discount, because the discount is already permanently applied.
- The account's own history shows a dedicated Mother's Day campaign in a previous year, now paused. Its learnings are worth extracting before rebuilding.

19. Projected outcomes, timelines and risks

These projections assume the tracking remediation in section 11 completes in week one and that the checkout tests in section 7 are implemented. If the checkout work does not happen, the ROAS projections below do not hold, and we would rather say that now than explain it in month three.

WINDOW	EXPECTED STATE	RECONCILED ROAS
Days 1-14	Tracking fixed. Reported ROAS likely falls as duplicate events are removed. This is correction, not decline.	True baseline established
Days 15-30	Checkout tests live. Retargeting restarted. Cart-to-checkout begins moving.	Baseline to 2.4
Days 31-60	Creative engine producing. Cart-to-checkout 30%+. First scaling step.	2.4 - 2.9
Days 61-90	Consolidated structure, 8+ live concepts, retargeting mature.	3.0+

What will look worse before it looks better

Two things, both expected, both worth stating in advance so they are not misread as failure.

Reported ROAS will drop in week one. Removing duplicate purchase events removes revenue from Meta's reporting that was never real. The number on the dashboard will fall while the business is unchanged. This is the point at which an agency without a written prediction gets fired for fixing something.

Consolidating campaigns and migrating to value optimisation resets the learning phase. Expect seven to ten days of unstable delivery and elevated cost per purchase after each structural change. This is the cost of the change, and it is why the plan sequences changes rather than making them all at once.

Risks

RISK	LIKELIHOOD	MITIGATION
Checkout changes are not implemented	Medium	The largest single risk in the plan. Tests 2, 3, 6 and 7 are deliberately low-effort so progress is possible even without theme development resource.
True ROAS is materially below 2.17	High	Expected. Targets re-baselined in week one. The improvement percentages hold even if the starting point moves.
Margin does not support 3x	Unknown	Resolved by the margin figure. If 3x is thin, the target becomes contribution-margin ROAS and the ladder gates tighten.
Creative production cannot sustain cadence	Low	Production sits with us. The only external dependency is one photography session.
Q4 CPM inflation compresses efficiency	High	Seasonally adjusted gates. Slate loaded and tested before 1 November.

20. Reporting and working cadence

Weekly

- Monday scorecard: spend, reconciled ROAS, purchases, CPA, AOV, cart-to-checkout rate, and creative-level performance. One page.
- What changed in the account last week, what is changing this week, and why.
- Creative pipeline status: in brief, in production, in test, scaled, retired.

Fortnightly

- Creative review. Assets reviewed individually, not in batches.
- Test readout: what was tested, what the result was, what happens next.

Monthly

- Full performance review reconciled against Shopify rather than against Meta reporting.
- Scaling gate assessment against section 15.
- Competitive re-sweep of the Ad Library. This category moves, and the positioning gaps in section 13 will not stay open indefinitely.

Always

- Cart-to-checkout rate reported every week until it clears 40%. It is the number that determines whether this account reaches 3x, and it should be visible in every conversation until it is fixed.
- Direct access. No account manager relaying messages to a media buyer.

21. Summary of findings

The account is not the problem. Meta scores this account 91 out of 100 on structure and delivery. CPMs are competitive, frequency is healthy, and the primary prospecting campaign returns 2.49. The media buying has been done reasonably well.

- 1,964 people added to cart in 90 days. 390 reached checkout. 192 bought. Cart-to-checkout runs at 19.9% against a 50-65% norm, and confirms independently in clean site-wide data at 16.4%.
- Taking that single step from 19.9% to 30% delivers the 3x ROAS goal on existing spend, existing creative and existing targeting. Taking it to 40% delivers 4.36.
- The purchase pixel is firing twice and the Conversions API has been offline for four days. The reported 2.17 ROAS is an upper bound and must be reconciled against Shopify before any target is set or any budget is moved.
- Meta is currently claiming credit for roughly 85% of all store orders, including organic and returning customers. That is beyond normal attribution overlap.
- Retargeting is switched off. 1,574 people abandoned carts in the window and are receiving nothing, against a 36,000-customer base with a natural repurchase cycle.
- Nearly every ad set optimises for purchase count rather than value, on a catalogue spanning \$120 to \$550. Meta flags this on eight separate campaigns.
- Three assets carry 48% of spend. One under-funded asset returned 7.46 ROAS at a \$47.71 cost per purchase on 2.4% of the budget.
- The category's advertising is 422 active US ads of keyword-stuffed product titles. The birth-statistics territory this brand actually owns returns 8 active ads, and effectively zero real competitors.
- Product photography on the live site is AI-generated with the original export filenames intact, on a brand whose entire proposition is that the object is real and handmade.
- Every SKU is permanently 20% off, which removes the promotional lever before Q4 and Mother's Day and compresses margin on every order.

Where this leaves us

Fix the measurement so decisions run on real numbers. Fix the one checkout step that loses four out of five buyers. Restart retargeting against a customer base that is currently being ignored. Then scale, one gated step at a time, against creative territory that nobody else in the category has claimed.

The 3x goal does not require a bigger budget. It requires the traffic you are already paying for to be able to finish buying.