

CONVERSION TRACKING

PROFIT REPORTING

MARKETING MIX MODELING

Their healthiest looking metric was **the warning sign**

A Shopify brand followed a record January with a heavy loss in February. The returning customer rate everyone was proud of was the clue, and the conversion data was too inflated for anyone to see it. We rebuilt the measurement, moved the budget onto new customers, then modelled where it should go next.

A US direct to consumer brand on Shopify Enthusiast category · name withheld at the client's request

3 to 6×

More conversions the ad platforms were reporting than were actually happening

95%

Tracking accuracy after the rebuild, measured against the store's own orders

7×+

Return on the two channels the model showed were carrying the business

01

The problem

02

Fix the data

03

Change the target

04

Prove it with a model

05

What it produced

INDUSTRY

Direct to
consumer
ecommerce

MODEL

Shopify, single
storefront

STACK

Shopify · GA4 · GTM web and
server · Google Ads · Meta ·
Bing · Klaviyo · SMS · Looker
Studio · Google Meridian

ENGAGEMENT

Measurement rebuild, then an
ongoing advisory retainer


PERIOD

January 2025
onward

A record month, then a loss, and nothing in the reporting explained why

January was the best month the business had seen. The team had gone all in on offers to their existing customer base and it worked. Four weeks later the same business posted a heavy loss. Nothing had broken. The offers had simply run out of people to land on, and no report in the company was built to show that.

Monthly net profit, indexed to January



Monthly net profit, indexed to January = 100. February reversed into loss. Absolute figures are the client's and are not published.

1

The platforms were reporting three to six times too many conversions

Google Ads and Meta were both counting far more conversions than the store recorded. Every cost per acquisition, every ROAS, every decision to scale or cut a campaign was being made against numbers that were wrong by a multiple nobody had quantified.

2

The budget was landing on people who had already bought


Retargeting the existing base is the cheapest looking spend in any account, because those people convert. It flatters every metric while quietly doing none of the work that keeps a business growing.

3

Nobody was looking at profit

Leadership had revenue dashboards. Revenue was fine in February. Profit was not, and there was no view in the business that put spend, margin and sales in the same place, so the reversal was only visible once the accounts closed.

The metric that told the real story



A returning customer rate this high is usually reported as loyalty. Read next to flat revenue, it says something different: the brand had stopped reaching anyone new, and was recycling the same audience at increasing cost.

Sound familiar?

- Your retention numbers look excellent and your growth has quietly stopped.
- Platform-reported conversions and your store's order count do not match, and nobody knows which to trust.
- Your dashboards show revenue. Nobody can pull up profit without building a spreadsheet.
- Promotions get decided in the week they run, rather than planned against a target.

02 · FIX THE DATA

Nothing else is worth doing until the numbers are true

There is no point arguing about budget allocation when the inputs are inflated three to six times over. The first job was to make the reporting match reality, and to put it somewhere leadership would actually look.

BUILT

- **Conversion tracking rebuilt from scratch** in Google Tag Manager, with server-side delivery where the platform allowed it, so purchases reach the ad platforms reliably rather than depending on what survives the browser.
- **Reconciled against the store's own orders** until tracking accuracy reached 95%. That number is the point of the whole exercise: it is the difference between a reported conversion and a real one.
- **A full UTM convention** applied across every platform, so campaign-level performance can be compared on one consistent basis instead of five different ones.
- **Profitability dashboards in Looker Studio**, built for leadership rather than for media buyers. Net profit next to revenue, spend and margin, in one view, updated daily.
- **Live campaign alerts into Slack**, so the media buyers see performance move during the day rather than reading about it a week later. Shortening that loop changes what people can actually act on.

Why the accuracy number matters more than it sounds

An account optimising against inflated conversions does not fail loudly. It quietly learns the wrong lesson: it bids up the audiences that appear to convert, which are usually the people who were going to buy anyway. Fixing the count is what lets the bidding algorithms start working on the real problem.

Why profit, not revenue

February had perfectly acceptable revenue. A revenue dashboard would have shown a decent month. The reversal only exists once you put spend and margin next to sales, which is exactly why that view had to be built before any strategy conversation was worth having.

03 · CHANGE THE TARGET

Stop paying to reach the customers you already have

With trustworthy numbers in place, the pattern was obvious and the fix was uncomfortable. A large share of spend was going to audiences the business had already acquired. That spend looks efficient on a platform report and does almost nothing for growth.

CHANGED

- **Budget moved off repeat-audience retargeting** and onto prospecting, with existing customers excluded from acquisition campaigns rather than quietly counted as wins.
- **Lead generation funnels introduced**, so paid traffic had somewhere to go other than straight to a product page. Not every visitor is ready to buy today, and the previous setup had no way to keep the ones who were not.
- **Nurture rebuilt across email and SMS**, so the list became a channel that develops a relationship rather than a discount broadcast.
- **Promotions planned against a calendar and a margin target** rather than decided reactively. Discounting without a plan is how a strong January becomes a difficult February.

The counterintuitive part

The returning customer rate went down, and that was the goal. When a business is genuinely acquiring, the share of orders from existing customers falls, because the denominator is growing. A metric that looks like it is deteriorating can be the clearest evidence the strategy is working. Agreeing that in advance, with the team, is what stops it being read as a problem three months later.

A year of weekly data, and a clear answer on where the next dollar goes

Once the data was reliable, we built a Bayesian marketing mix model in Google Meridian on 52 weeks of weekly data across five paid channels, with controls for seasonality and promotions. The point of a model is not to replace judgement, it is to answer the one question platform attribution structurally cannot: which spend is actually incremental.


How it works, briefly

Each channel's spend passes through two transformations before it reaches the revenue equation. **Adstock** carries the effect of a campaign forward into the weeks after it ran. **Saturation** captures the fact that the tenth dollar earns less than the first. Every parameter is estimated as a distribution rather than a single number, which is what produces honest uncertainty on the outputs.

Whether to trust it

The model explained 89% of weekly revenue variation, with an 11% error weighted by revenue, meaning it is most accurate in the high-revenue weeks where budget decisions actually get made. Strong enough to set allocation. Not precise enough to forecast a specific week, and we say so.

Average return versus return on the next dollar



Where the two bars are far apart, the channel has been a good investment historically but has little room left to grow.

Average return grades the spend already made. Marginal return is the one that should set next quarter's budget. Two channels cleared 7x on average, and one of them still had room to grow.

FINDING

Two channels were carrying the business, and the model was confident about both

Paid search and email each returned more than 7x on average, with tight credible intervals. The three remaining channels sat between 1.6x and 2.2x with wide intervals, and their marginal return, the return on the next dollar, sat at or below break-even. Money moved into those channels was not buying incremental revenue.

FINDING

Historically excellent and still worth scaling are not the same thing

Email had the highest average return in the account and one of the lower marginal returns, which is the signature of a channel that is well run and close to its ceiling. Paid search had a slightly lower average and a clearly higher marginal return, meaning it was the channel with headroom. Reading only the average would have produced the wrong decision.

FINDING

There was very little demand underneath the advertising

The model estimated the baseline, the revenue that would arrive with no paid media at all, as close to nothing. For this business that is a strategic fact rather than a criticism: paid media is not amplifying existing demand, it is creating it. Continuity of spend is therefore not a lever to be switched off in a quiet month, and the brand-building work that would build a baseline is a separate, longer project.

The reallocation, and why no extra budget was needed

At the best possible allocation, the last dollar spent on each active channel returns the same amount. When one channel returns several times more on its next dollar than another, moving money between them increases total revenue without increasing total spend. That is the entire argument, and it is why the recommendation was a reallocation of roughly a fifth to a third of the weaker channels' budget rather than a request for more.

Trustworthy numbers, a growing customer base, and profit back at a record

BEFORE	AFTER
Platforms reporting three to six times more conversions than happened	Tracking reconciled to the store's own orders at 95% accuracy
Five platforms, five campaign naming conventions	One UTM convention, so campaigns are comparable across channels
Revenue dashboards only, profit visible after the accounts closed	Daily profitability view built for leadership, with spend and margin alongside sales
Performance reviewed weekly, in arrears	Live campaign alerts, so buyers can act during the day
Spend concentrated on audiences already acquired	Acquisition prioritised, existing customers excluded from prospecting
Budget split by habit and platform-reported ROAS	Budget set by modelled marginal return, refreshed as the data grows

95%

Tracking accuracy against the store's own order data, from a starting point nobody could quantify

Reconciled directly against store orders.

0.84 → 0.76

Returning customer rate, falling deliberately as new buyers re-entered the mix

Monthly store data. A fall here is the intended outcome, not a decline.

136

July net profit, indexed to January = 100, on a smaller sales month than June

Client reporting. Absolute figures not published.

What we are, and aren't, claiming

A tracking rebuild does not create demand and a model does not spend money well on its own. What changed here is that the team could see what was actually happening, in time to act on it, and had an independent answer on where the next dollar should go. **The business did not need more budget or more tools. It needed the numbers to be true and the target to be the right one.** The profit recovery followed the decisions that became possible once both were in place.

Method notes. Tracking rebuilt in Google Tag Manager with server-side delivery where supported, reconciled against store order data until accuracy held at 95%. Dashboards in Looker Studio across the store, ad platforms and analytics. Marketing mix model built in Google Meridian: Bayesian regression, weekly grain, 52 weeks, five paid channels, geometric adstock and Hill saturation, controls for seasonality and promotional periods, Hamiltonian Monte Carlo sampling with credible intervals on every output. Fit checked before any figure was reported. Known limits stated in the client's report: 52 weeks is a short window for saturation shape, correlated spend across channels in peak weeks limits separability, and parameters are treated as stable over time, so the model is re-run rather than trusted indefinitely.

Is your best metric telling you what you think it is?

High retention with flat growth, platform conversions that do not match your store, dashboards that show revenue but not profit. Each one is worth checking on its own. Together they are a pattern, and it is a fixable one.

[Book a measurement review](#)

Want a cheaper starting point? [Run a free Google Ads audit](#) →

About Swishiy

Swishiy is a data and measurement agency for ecommerce and B2B brands that spend real money on paid media. We sit between the platforms that report on themselves and the agencies that report on the systems they manage: the independent layer that verifies whether the numbers are true, fixes them when they aren't, and builds the measurement that lets a business bid on value rather than volume.

Tracking & conversion audits

GA4, GTM, consent, Google Ads, Meta, Microsoft. Every defect with its evidence and its commercial consequence.

Server-side & CRM pipelines

First-party server containers you own. Offline and CRM conversions delivered to every platform with value attached.

Marketing Mix Modeling

Bayesian MMM (Google Meridian) to show which channels drive incremental revenue, independent of platform attribution.

Track record

Top Rated Plus on Upwork · 77+ projects · 100%
Job Success. Built and led by Qudees Kiani.

Contact

swishiy.com · qudees@swishiy.com
Book a call: swishiy.com/book



swishiy.com · Trust-Driven Growth · Case Study No. 04