

MARKETING MIX MODELING

DASHBOARD

SERVER-SIDE TRACKING

Judge the ads on new customers, not on everyone

A sports memorabilia brand with a loyal customer base was reading its ad performance from six places that never quite agreed. We built one dashboard, tracking that holds up, and a marketing mix model, then pointed the same budget at the customers they did not have yet.

Powers Sports Memorabilia Autographed memorabilia & signings · United States ·

powerssportsmemorabilia.com

2 years

Of weekly data modelled, and run against two different definitions of success

+50%

More new customers per day on days the ads ran versus days they were paused

Same budget

The reallocation needed no increase in spend, only a change of target

01

The problem

02

One view of the truth

03

What the model said

04

What changed

05

The retainer

INDUSTRY

Sports memorabilia & autograph signings

MODEL

DTC ecommerce, Shopify

STACK

Shopify · GA4 · GTM web + server (Stape) · Google Ads · Meta · Klaviyo · Looker Studio · Google Meridian

ENGAGEMENT

Analytics & paid-media retainer, ongoing

PERIOD

Mar 2026 – present

Six sources of truth, and no single place to compare them

Powers Sports Memorabilia has been selling authenticated autographs and running athlete signings for years. That builds something most ecommerce brands don't have: a large base of collectors who come back on their own, and a brand people search for by name. It also creates a specific measurement trap.

1

Platforms grade their own homework

Every ad platform reports on itself, counts on its own attribution window, and has no view of the others. When a brand is well known enough that people search for it by name, some of what the platforms claim is demand that already existed. That is true of every account, and no platform will tell you which part.

2

The picture was split across six tabs

Sales in one system, spend in two ad platforms, behaviour in analytics, email in another tool. All of it accurate, none of it side by side, and every source counting on a different window. Comparing them by hand each week is a job nobody should have to do.

3

New customers were the blind spot

No ad platform reports on whether a sale came from someone new or someone returning, yet for a brand with a loyal base that is the only question worth asking of the ad budget. Without it, you cannot tell acquisition from harvest.

Sound familiar?

- Your brand is strong enough that people search for you by name, and no platform separates that from genuine acquisition.
- You look at ROAS in Google, ROAS in Meta and revenue in your store, and the three never quite agree.
- You can't say how many of last month's customers were new, or which channel brought them.
- Someone has suggested spending more, and you have no independent way to check whether that is right.







What we set out to do

Three things, in order. Build one view of the business that puts Shopify, the ad platforms and analytics on the same page. Make sure the data feeding that view is real. Then use a marketing mix model, not platform attribution, to decide where the money should go. The engagement has since become a monthly retainer, which section 05 describes.

A dashboard the owner actually opens, built on tracking that holds up

The dashboard came first, because it changes the conversation. Once yesterday's sales by channel sit next to what each platform spent, new versus returning, the question stops being "what does Meta say?" and becomes "what did we get for it?"

The Looker Studio dashboard, six views



Shopify, Google Ads, Meta, GA4 and Klaviyo blended into one report. Every page filters by date, channel, region and customer type. Shown as a wireframe; the live report carries the client's figures.

What's in it

- Daily sales, orders, AOV, blended ad spend and ROAS, with period-over-period change
- Yesterday's paid performance: sales by channel, top products, new vs. returning
- GA4 by source and medium: users, revenue, conversion rate, engagement
- Conversion funnels from first visit to purchase, and the user journey by channel
- Meta and Google Ads drill-downs to campaign, ad set, ad and keyword

Tracking that holds up

A dashboard is only as good as what feeds it, so we rebuilt the measurement underneath.

- **Server-side tracking** on a first-party Stape container, so purchases reach GA4, Google Ads and Meta reliably despite ad blockers and Safari's cookie limits
- **A platform quirk found and filtered:** a store-platform sandbox frame was being logged by analytics as if it were a landing page, which inflates reported traffic and drags every conversion rate down with it. Not a setup error, and easy to miss until someone goes looking
- **Meta pixel replaced** and the Conversions API set up, with a new-customer exclusion list synced from Shopify
- **Consistent UTMs** across Google, Meta, email and SMS so Shopify and GA4 tell the same story

Why the tracking work matters to the model

A marketing mix model learns from two years of weekly spend and sales. If the sales data is inflated or the channel labels are inconsistent, the model learns the wrong lesson with full confidence. Fixing the plumbing first is not a nice-to-have; it is what makes the next section trustworthy.

Two years of data, two definitions of success, two different answers

We built a Bayesian marketing mix model in Google Meridian on roughly two years of weekly data: paid spend by channel, email sends, organic and social traffic, search impressions, and holiday flags. It accounts for carry-over (an ad keeps working after it runs) and diminishing returns (the tenth dollar earns less than the first). Then we ran it twice: once against total sales, once against new-customer sales only.


Where revenue comes from, by lens



The same two years of data, scored two ways. Against total sales, most revenue comes from baseline demand. Against new customers only, paid media and email do about half the work. Shown as relative shares; exact figures are not published.

That split is the whole story, and it is true of any brand with real repeat custom. Score the ads against total sales and they look like a modest contributor to a business that largely runs itself. Score them against new customers only and they are the acquisition engine. Both readings are correct. The one that should drive the budget is the second, because acquisition is the only thing the ads can actually do.

Return per dollar by channel, with uncertainty



Relative return per dollar on the total-sales lens, scaled to break-even. Whiskers are the model's credible interval, which is the honest way to present a model. A narrow interval well clear of break-even is a decision. Grey bars are the ones a single lens cannot settle, and those are tests to run rather than verdicts to act on.

FINDING

Brand demand was flattering the campaigns meant to find new customers

Performance Max ships without brand exclusions by default, so it will happily serve people who already typed the brand name. Their reported return looks strong and their real acquisition work is hidden underneath it. This is the single most common thing we find in well-known brands' accounts, and it is a settings change rather than a spend decision. It had to come first, because until it is fixed every other number is measured against an inflated one.

FINDING

Email returned several times more per dollar than any paid channel

Klaviyo and SMS were the most efficient revenue in the account by a wide margin, at close to zero marginal cost. The list had been under-used. Growing it and sending to it more was the cheapest revenue available.

FINDING

The two lenses disagreed about Meta, so we did not let one of them decide

On the total-sales lens Meta sat around break-even, which is exactly what you would expect from a prospecting channel measured against a business carried by returning customers. The easy call was to cut it. We did not make that call, because the new-customer lens pointed the other way and there was a third way to check.

The natural experiment: new customers when Meta ads were on vs. off



Average new customers per day, indexed to paused days = 100. Five separate pause windows over roughly 16 months acted as a natural experiment; the gap is statistically significant. Same business, same period, only the ads switched off.

The account had been fully paused five separate times over the period, for unrelated reasons. Comparing those days to the rest isolates the ads' effect from seasonality and trend. Same business, same period, only the ads switched off.

One lens is never enough

Platform attribution was enthusiastic. The total-sales model was lukewarm. The new-customer model and the on and off comparison both said the channel was bringing in buyers who were not otherwise coming, at a rate that clears the cost of media and fees once you price in what a new collector is worth over the following years. The right answer was neither "more Meta" nor "no Meta". It was to stop both Meta and Google bidding on people already on their way in, point them at new customers only, and judge them on new-customer value from then on. Three methods, one decision, and it would have been the wrong decision on any one of them alone.

Same budget, different target

No increase in spend was required. The reallocation moved money away from where the platforms were claiming credit and toward where the model showed incremental new-customer revenue, then set up the experiments to keep checking.

GOOGLE ADS

- **Brand exclusions added to Performance Max**, so branded demand is served by the branded campaign and PMax has to earn its return on new customers.
- **Non-brand search rebuilt** around the terms already converting well at small spend, with budget moved off the terms that were not earning it.
- **Search-term review every month**, splitting brand from non-brand so the true cost of acquiring a new customer is visible rather than blended into a number that brand demand is flattering.

META

- **Rebuilt for new-customer acquisition**: an Advantage+ shopping campaign with a new-customer budget cap, past purchasers excluded from all prospecting, retargeting kept small and separate.
- **Campaigns built around what the brand is known for**, signings and marquee athletes, with an audience architecture (product viewers by sport, high-value purchaser lookalikes, signing-page visitors) documented and handed over.
- **A clean pixel and CAPI** so Meta's own reporting stops disagreeing with Shopify.

EMAIL AND THE YARDSTICK

- **Email and SMS re-prioritised** as the highest-return channel in the account, with list growth promoted to a paid-media objective in its own right.
- **Customer lifetime value modelled** from six years of customer history, by acquisition cohort, so the owner knows what a new customer is worth in year one and beyond.
- **A break-even ROAS the account has to clear**, including media and fees, at each level of spend. Every platform's reported ROAS is now read against that number rather than against nothing.

One thing we recommended, tested, and killed

We proposed a lead-capture landing page to grow the list from paid traffic. It captured leads quickly and it looked like it was working. When we matched those leads against actual orders, almost everyone who went on to buy turned out to be an existing customer being re-identified through the form rather than a new one being acquired. We showed the data and recommended stopping it. Checking your own ideas as hard as you check everyone else's is the part of this job that makes the rest of it worth trusting.

What the business has now

One report to open in the morning. A model that says which channels earn their keep, with honest uncertainty attached. A yardstick every spend decision is measured against. And ads aimed at the customers they do not have yet rather than the ones they already do.

What we are not claiming

A mix model is a decision tool, not a crystal ball. It is reliable on direction, less so on precise dollar forecasts, and we say so in every report. The model is re-run as the reallocation beds in, so the next set of decisions is made on the results of the last set.

05 · THE RETAINER

What a month looks like

The project work above turned into an ongoing engagement. This is what lands each month, and it is what any retainer client of ours gets: the analysis, the decisions and the plumbing, from one person who has seen all of it.

DELIVERABLE	CADENCE	WHAT IT ANSWERS
Live performance dashboard	Always on	What happened yesterday, by channel, product and customer type. Shopify, Google, Meta, GA4 and Klaviyo in one place.
Performance diagnostic	Monthly	New-customer acquisition, paid-media efficiency and site conversion, year over year, with the five things that matter and what to do about them.
Google Ads structure & search-term review	Monthly	Brand versus non-brand spend and return, terms to scale, terms to retire, campaign structure changes.
Meta campaign plan & audience architecture	Per launch	Campaigns, ad sets, audiences, exclusions and creative direction, built around new-customer acquisition and the signings calendar.
Customer value & break-even analysis	Quarterly	Customer value by cohort, cost to acquire, and the ROAS the account must clear at each spend level after fees.
Marketing mix model refresh	Quarterly	Re-run on the latest data to confirm the reallocation worked and set the next one.
Experiment design & read-out	As needed	Lead-capture tests, ads on/off analysis, campaign experiments, each with a clear verdict.
Tracking maintenance	Ongoing	Server-side container, GA4, pixels and UTMs kept clean. Defects found before they reach a decision.

Methods, briefly. Looker Studio on Shopify, Google Ads, Meta, GA4 and Klaviyo connectors. Server-side GTM on Stape. Google Meridian Bayesian MMM, weekly grain, two years of data, adstock and Hill saturation, two KPI runs (total sales and new-customer sales), convergence checked before any number was reported. Ads on/off comparison using Welch's t-test across five pause

windows. LTV from six annual customer snapshots. All figures in this study are rounded or indexed at the client's request; the client's reports carry the exact numbers.

Are your ads finding new customers, or meeting the ones you already had?

If your brand has real organic pull, no ad platform can separate the demand it created from the demand it merely met. A marketing mix model is the independent check. It starts with a scoping call to see whether you have the data history to support one.

[Book an MMM scoping call](#)

Not sure your tracking is clean enough to model? [Start with 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

