

 CASE STUDY · MEDIA MANAGEMENT

From booked rides to **real margin**

How we realigned Allocab's Google Ads bidding, from conversion volume to the business value of each ride.

CLIENT

allocab

Ride-hailing · Mobility
French market

AUTHOR

Léo Llinares

DATE

June 2026

READING

~7 minutes

SCOPE

Google Ads · Search & PMax



Communication rule. Only margin indicators are quantified (as % of change or € earned per ride). No absolute ROAS, CPA or volume figure is published.

— IN SHORT

Key figures

On a mature acquisition account, switching from booked-ride bidding to real-margin bidding produced a double effect, isolated in an AB test and then confirmed at the scale of the whole setup.

+42%

margin per completed ride

8-week AB test, test campaigns (Margin tROAS) vs control (tCPA): +42% margin per completed ride and +19% total margin, while spending less than the control. Order of magnitude: the extra margin on each ride represents about a quarter of what acquiring a ride used to cost before the switch.

+26%

margin generated on Google Ads non-brand Search and Performance Max

First full month after the switch, at near-stable investment (+2.8% cost) vs the previous month. With no volume-for-value trade-off.


01 — FOREWORD

An ever more contested paid market

Allocab operates in a French ride-hailing market dominated by international platforms that heavily over-invest in paid acquisition. Competitive pressure on Search is structural and growing: the category's advertising investments in France have increased year after year since 2022, and Allocab's competitors spend substantially more than it does on the channel. The French ride-hailing market keeps growing in ride volume (+16% in 2024 according to ARPE), but for a challenger, every position on Search costs more to defend each year.

For a player whose acquisition still hinges on Search, the challenge is twofold: holding ground against competitors who keep outspending, and making every euro invested produce more value. When you cannot win the volume war, you win the value-per-conversion war. This is precisely the trajectory we have been driving with Allocab since late 2024.

This note covers a structural switch made in spring 2026: moving from bidding on booked rides to bidding on the real margin of each completed ride. Method, results, and what we take away for other mature acquisition accounts.



02 — THE STARTING POINT

Booked-ride bidding was reaching its limit

Until early 2026, Allocab was bidding on the **booked ride**: the most natural conversion event to track, the highest-volume one, the one that gives the Smart Bidding algorithm the most signal to learn from.

But bidding on booked rides means bidding on an intermediate event, not on real business value. The account had reached the limit of what that signal could deliver, for three operational reasons.

- 1 Not every booked ride is completed.** The completion rate varies with context (weather, driver availability, customer behaviour). At an identical target, the algorithm cannot tell a ride that will be honoured from one that will not.
- 2 Unit margin varies widely from one ride to another.** Depending on customer profile (consumer, SMB, key account), trip type (short urban, airport, long distance) and vehicle type. In Q1 2026, margin per ride across campaigns varied by up to a factor of 1.8 between the least and the most profitable.

3 The algorithm was arbitrating blind to those gaps. Smart Bidding was optimizing the volume of booked rides at a target cost, blind to the real profitability of each trip. The resulting mix was optimal on booked rides, not on margin.



OUR TAKE

Bidding on booked rides is a good first maturity stage. It reaches its limit when you want to improve the quality of your conversion mix, not just its unit cost. At that point, **changing the event the algorithm optimizes is more effective than endlessly refining targeting** on an event that only imperfectly reflects business value.

03 — THE APPROACH

Three phases, eight weeks of testing before any rollout

We designed the switch in three phases, with a rigorous test protocol before any account-wide rollout. Methodological discipline is what made the effect measured and reproducible.

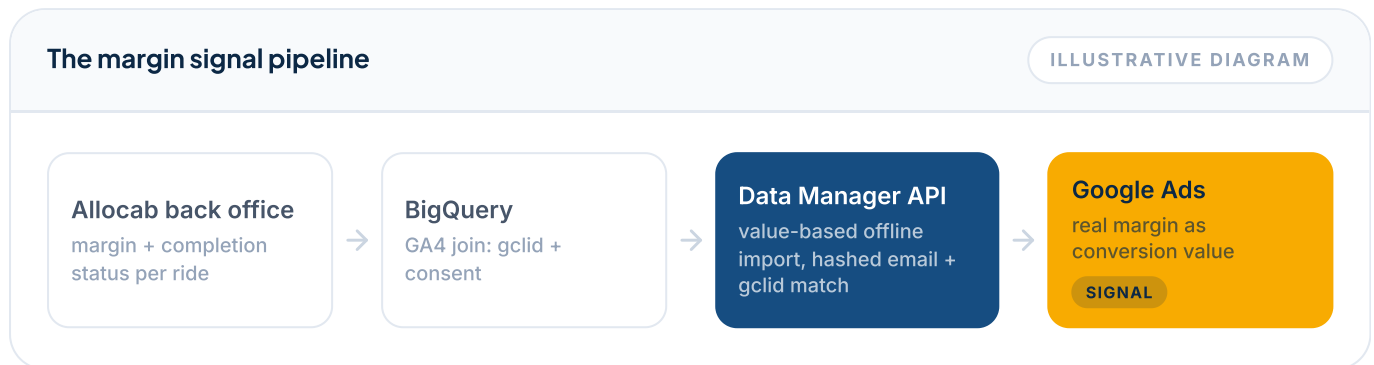
1 Step

Make the completed-ride and margin signal reliable

Four months of upstream preparation: sending, for each ride, the **margin** and the **completion status** from Allocab's back office to BigQuery, enriching that signal with Google Analytics 4 (gclid and consent status of each ride), then feeding it back into Google Ads through the Data Manager API as a value-based offline conversion import; the margin of each completed ride sent as conversion value, matched on hashed email and gclid. On the site side, server-side tagging and advanced Consent Mode secure a durable, consent-compliant measurement. The signal becomes usable as a conversion value by January 2026: Smart Bidding now has a measure of business value per conversion, rather than a mere counter of intermediate events.

This stage is more demanding than it looks. Bidding on completed rides means bidding on an event that happens later in the funnel, in lower volume, with extra loss points: rides not honoured, CRM data not sent back in time, attribution gaps

between the booked and the completed event. The error tolerance is very low: any hole in the feed immediately degrades learning.



ABSOLUTE PREREQUISITE

Without rock-solid tracking of completion status and a clean, historized margin signal, no Margin tROAS bidding is possible. **This stage often weighs more in the outcome than the technical switch itself.**


2 Step

Test before switching: an eight-week AB test

Controlled test launched on 10 February 2026 on four mirrored campaigns: **two test campaigns** duplicated on Margin tROAS bidding, **two control campaigns** kept on the existing tCPA. Three successive measurement windows over eight weeks, with cross-checked sources: Google Analytics 4 (volume and cost) and Allocab's CRM (real margin and completed rides). At each window, we read the gap across four dimensions: volume, acquisition cost, unit margin and total margin.

Margin per completed ride, test vs control

ILLUSTRATIVE DATA · BASE 100



↗ **Window 3 verdict: +42% margin per ride, while spending less than the control**

Illustrative. What matters is not the absolute level of the first windows but the dynamic: the gap narrows week after week until the test pulls ahead.

Across the first two windows, the test lagged the control on almost every metric in absolute terms. This is where discipline matters most: not concluding too early. When you ask the algorithm to optimize an event it has never targeted, the learning phase is mechanically long.

What we observed over the eight weeks of testing:

- ✓ **Windows 1 and 2.** The control stays ahead on most indicators, but the test's trajectory shows clear learning, with the gap narrowing at every measurement.
- ✓ **Stabilization.** After a few weeks, parity between control and test on the key indicators.
- ✓ **Window 3 (verdict).** The test pulls clearly ahead. At comparable spend across the four aggregated campaigns, +42% margin per completed ride and +19% total margin, while spending less than the control.

3 Step

Switch: consolidation and migration in waves

On the strength of the test verdict, the full switch was carried out in April 2026. **The switch was also an opportunity to simplify the account structure.** With a unified margin objective, the case for multiplying campaigns by segment type collapses: margin being intrinsically comparable, the target is unified across the account. Everything can be consolidated into large Search and PMax campaigns,

with segmentation at ad-group level, which mechanically simplifies the algorithm's learning.

In practice, the migration followed two principles.



A unified target tROAS across prospecting, kept fixed during the migration phase so as not to disturb the algorithm's learning.



Migration in weekly waves. The AB test campaigns, already run in, served as the foundation. The rest of the account was added progressively, ad group after ad group. Switch completed at the end of April.



WHAT WE DELIBERATELY RULED OUT

Switching the whole account at once. It is tempting to save time, but it creates a simultaneous learning gap across all campaigns, precisely when the algorithm needs to find its bearings. **Sequencing in waves smooths the learning loss** and lets you read the first signals before propagating.

Expected and observed effect: the later waves learned much faster than the initial AB test. Where the test had needed nearly eight weeks to pull clearly ahead of the control, the switch waves on the rest of the account stabilized within two to three weeks. The reason is mechanical: the account had by then accumulated enough historical data on the downstream event (completed ride and its margin) for the algorithm to no longer start from scratch at each wave. The upstream signal-reliability work pays off a second time here, by speeding up subsequent switches.

04 — THE RESULTS

Four dimensions improving together

At near-stable investment vs the previous month, the first full month after the switch (May 2026) delivered **+26% margin** on Google Ads non-brand Search and Performance Max. Beyond the figure, the striking point is the absence of any volume-for-value trade-off.



Acquisition quality and **conversion volume** up: more new customers, more completed rides.



Cost efficiency and **mix value** up: a lower unit cost, on a more profitable mix of rides.

That is a rare configuration in mature paid acquisition, and it is exactly what margin bidding is meant to produce: aligning the algorithm's optimization objective with the real business value of each conversion, without sacrificing volume.



WHAT THE EIGHT WEEKS OF TESTING REVEALED

The advantage is not a one-off effect: it grows as the model learns. The AB test compared both methods in parallel, same campaigns, same market conditions. It is **causal evidence, isolated from seasonality, price and scaling effects**, and the gap widened window after window once the learning phase was over. The full April switch replicated the effect at the scale of the entire setup, where it kept strengthening as history accumulated on the downstream event. Both readings converge on the same signal: a gain that consolidates over time.

In a market where competitors keep investing more in paid, this quality of unit-level arbitration becomes a key competitive advantage, with no volume trade-off.

05 — EDGEANGEL'S TAKE

When to make the switch

Real-margin bidding is not the right entry point for every account. To get full value from it, three conditions must be met.

1

An account that has matured on its intermediate event. As long as bidding on the upstream event (booked ride, qualified lead, add-to-cart, sign-up) has not stabilized, switching to a downstream event is premature. The intermediate event is the algorithm's school; real business value is its specialization.

2 A reliable business-value signal that can be fed back per conversion.

Margin, lifetime value, unit contribution, effective service rate. Without that signal, the algorithm has nothing to optimize beyond the cost of an intermediate event. Making it reliable is a data project in its own right, and it is what determines success, not the technical switch.

3 Enough conversion volume for learning. Margin tROAS asks the algorithm to arbitrate between conversions, which requires dozens of conversions per week. A downstream event is, by construction, lower-volume than an upstream one. On a low-volume account, variance drowns the signal.



OUR TAKE

When these conditions are met, switching to real margin takes the account past a performance ceiling that cannot be reached on the intermediate event. When they are not, the effort is better spent building them. **The full chain, from feeding the signal back to switching the campaigns, is a long-haul project: the right maturity milestone for a paid acquisition account in 2026.**

06 — METHODOLOGICAL SCOPE

What is measured, and what is not

Data sources	Google Analytics 4 for consolidated acquisition volumes and costs, Allocab's CRM for real margin and completion status. Google Ads platform signals are tracked as a complement, to be read with caution as they are not consolidated with the business view. Cross-referencing done in BigQuery on a daily window.
Case scope	Google Ads non-brand Search and PMax exclusively (excluding App campaigns, Demand Gen and Microsoft Ads).
AB test measurement	Synchronous test vs control, same duplicated campaigns, same market conditions. Three successive windows over eight weeks. The Window 3 verdict, the latest and most stable measurement, is the one used for the published figures.

Post-switch month

MoM comparison April 2026 vs May 2026, May being the first full month in stabilized regime after the switch was completed at the end of April. Near-stable investment (+2.8%). Figure taken from the Google Ads platform (conversion value of the margin action, non-brand Search + Performance Max scope, excluding brand, App and Demand Gen), reproducible in the account. To be nuanced: a MoM comparison includes seasonality, so the AB test remains the causal evidence. The YoY comparison provides a complementary reading at comparable seasonality.

— TO GO FURTHER**Causmos & Causal Impact: measuring Social Ads campaign effectiveness**

The same methodological rigour applied to another channel

**Google Analytics: data modeling and attribution**

The attribution foundations that make this kind of switch possible

**A closer look at cross-channel budgeting in Google Analytics**

Tomorrow's media management

LET'S TALK

Facing a similar challenge on your acquisition account?

Real-margin bidding is a maturity milestone that does not deploy in a few weeks. If you are wondering whether your account is ready, or how to build the data chain that would make it possible, we are happy to talk it through.

[**Contact us →**](#)