A PRACTICAL GUIDE FOR BRANDS, RETAILERS AND ANALYSTS

Make the promotion earn its place.

Three reproducible workflows for contribution, experiment design and offer allocation. No opaque scores. No invented commercial results.

A larger sales number is only the beginning.

A discount can raise volume while reducing contribution. A precise estimate can still describe a biased comparison. An attractive offer can consume the budget that would support a better combination.

Decision	Useful question	Start with
Review	What did we give up to gain volume?	Campaign economics
Experiment	What could this design teach us?	Independent units
Plan	Which feasible offers work together?	Candidate policies

Use the comparison that matches the decision.

The following cases use fictional inputs and the same calculations as MarginTrial. They show how to inspect a result, not what any merchant achieved. Inputs, assumptions and omitted effects belong beside every number.

Your source data stays on your device. The public previews work without an account. The downloadable editions run locally; the optional Agency batch runner needs Node.js 20 or later.

Open the three-workspace hub

This guide is independently prepared product education. It does not promise income, causal identification or a particular commercial outcome.

01 / REVIEW

Find the margin behind the sales lift.

Synthetic case: one product, the same seven-day window, one currency and no returns, funding or later displacement.

Input / outcome	No promotion	Promotion
Units	1,000	1,600
Unit price	EUR 50	EUR 40
Variable unit cost	EUR 25	EUR 25
Extra campaign spending	EUR 0	EUR 300
Revenue	EUR 50,000	EUR 64,000
Contribution	EUR 25,000	EUR 23,700

Revenue rises EUR 14,000. Contribution falls EUR 1,300. The offer gives away EUR 10 on every one of the 1,600 promotion units, including units that the supplied baseline says would otherwise have sold at the regular price.

Contribution bridge	Change
600 extra units at EUR 25 regular contribution	+EUR 15,000
EUR 10 price reduction on all 1,600 units	-EUR 16,000
Allocated campaign spending	-EUR 300
Total versus supplied baseline	-EUR 1,300

The bridge reconciles an accounting identity. It does not establish that the offer caused 600 extra sales. Review seasonality, availability, comparison periods and category effects before trusting the expected no-promotion units.

Reproduce this exact worked case

02 / EXPERIMENT

Measure at the assignment level.

The independent unit is the object assigned to treatment, such as a store. Repeated transactions within that store are not extra independent assignments.

Specify the design before observing outcomes.

Fix the eligible-unit list, allocation scheme and seed. Choose the primary outcome, worthwhile increase, measurement window, missingness rules and later-payback window. Keep allocation and analysis consistent: strata, matched pairs and unit changes describe different designs.

Fictional outcome summary	Control	Treatment
Independent units	24	24
Mean outcome	1,000	1,070
Sample standard deviation	120	135

Treatment minus control: +70. Welch standard error: 36.87; approximate degrees of freedom: 45.38. The two-sided 95% interval is **-4.24 to 144.24** and p is approximately 0.064.

The interval includes zero and a range of commercially different effects. It does not establish equivalence. Without a defended allocation design, report an association; precision cannot remove confounding. The software cannot verify a randomization declaration.

Plan for a useful detectable difference.

Under a normal approximation, SD 1, a target difference of 0.2, two-sided alpha 0.05, power 0.80 and equal allocation require about **393 independent units per arm** before losses. Clustering and attrition can raise the requirement. Use unit-level SD when the unit is already aggregated.

Open Experiment studio and its explicit assumptions

This is one comparison without repeated-peeking or multiplicity adjustment. A two-period DiD needs additional identification assumptions, including parallel untreated trends; one pre-period cannot establish them. Method references are on page 6.

03 / PLAN

Spend the budget on the combination.

Synthetic six-product case: 18 candidate offers, mutually exclusive within each product; a EUR 1,600 spending budget, 8,500 sold-unit capacity and a maximum 25% markdown.

Feasible policy	Contribution	Spending	Units
Regular price	EUR 49,100.00	EUR 0	4,460
Greedy benchmark	EUR 54,006.70	EUR 1,355	5,882
Proved best	EUR 55,812.60	EUR 1,570	6,179

The selected plan adds EUR 6,712.60 against the supplied regular-price baseline and leaves EUR 30 of the spending budget. Three deep offers are selected: Daily blend, Hydrating serum and Weekend candle. The other three products remain at regular price and still consume capacity.

A feasible plan and a proof are different outputs.

MarginTrial compares each product's candidate offers with a no-promotion policy. Bounded branch-and-bound searches combinations under stock, capacity, spending and markdown limits. If the search proves the upper bound equals the incumbent value, the result is optimal within the model.

A larger search may stop at its node or time limit. Read the remaining upper bound and optimality gap beside the feasible result. The search gap describes possible improvement under the model; it is not a confidence interval or uncertainty about demand.

Keep response assumptions visible.

Manual mode uses the expected volumes you supply. Elasticity mode is an explicit constant-elasticity scenario, not a fitted demand curve. Funding changes contribution but does not offset the campaign-spending constraint. Returns, substitution and future displacement are not inferred by this planner.

Explore the planner's reproducible six-product example

The example's optimum depends on its fictional inputs and cent-rounded policy objective. The public preview supports 8 products / 32 offers; offline Agency supports 120 products / 2,000 offers.

04 / KEEP THE EVIDENCE

Make the result reproducible.

A campaign table, an experiment table and a candidate-offer table contain different objects. Preserve the input grain as well as the final report.

Campaign review: one product per campaign row.

Required columns: **campaign, product, baseline_units, promo_units, regular_price, promo_price, unit_cost**. Optional fields describe allocated spend, support, returns, window lengths and later displacement. All rows share one currency and a consistent tax basis.

Experiment analysis: one independent unit or pair.

Independent outcomes use **unit_id, group, outcome**. Group summaries use **group, n, mean, sd**. Matched pairs use **pair_id, control, treatment**. Two-period DiD uses **unit_id, group, before, after**. The importer rejects invalid rows instead of silently retaining a partial sample.

Offer planning: one candidate policy per product row.

Start with product, offer, regular_price, promo_price, unit_cost and baseline_units. Manual response adds expected_units; elasticity mode adds its stated elasticity. Competing offers must share the same product baseline, unit cost, stock and capacity inputs.

Retain the portable files.

Export each workspace's own JSON backup, its relevant CSV and a readable report. Optional browser-local saving is convenient but is not cloud backup. A purchase-recovery file restores license access; it does not contain your analysis data.

Agency also includes a local Node batch runner. Its three commands - **audit**, **plan** and **experiment** - create JSON, CSV and standalone printable HTML without an API, npm dependency or network connection. The HTML previews at most 100 rows while totals use all valid input rows.

Read the exact schemas and workflow guide

Keep synthetic-data markers when working with the supplied examples. Editing fictional values does not independently verify commercial performance. Unmarked imports are supplied, unverified data.

05 / YOUR NEXT DECISION

Record what would change your mind.

Use this short record alongside the workspace exports. State the commercial threshold before treating a positive number as a decision.

Decision owner, horizon and intended action

Expected baseline or comparison design and its source

Costs, returns, funding and spillovers included or omitted

Worthwhile improvement and constraints

Result, uncertainty or remaining search bound

Evidence that would change the decision; review date

Use the free tools, synthetic cases and decision workspaces

Personal includes offline campaign review and Experiment studio. Agency adds portfolio planning, the local batch runner and client-use rights. Current availability and license terms are on the website; no subscription or external AI service is required.

Method references

R documentation: Welch and paired t procedures

J-PAL: power calculations and clustered designs

World Bank DIME: design, identification and data structure

References supply methodological context; they do not endorse MarginTrial or validate these fictional outcomes. All examples were independently prepared for this guide.

You may share this guide unchanged with attribution to MarginTrial. Software licenses are separate.