

AGENCIO · BERTHA · CLAUDE


The Microfactory — Advertising OS, as a skill

A full-stack marketing agency compressed into seventeen reference engines and one operating doctrine. Client materials go in; strategy, cohorts, creative, media plans, briefs, decks and reports come out — every output carrying its behavioural-science rationale. Reporting closes each cycle back into the next brief.

🔒 Private & proprietary — © 2026 Agencio. Copying and distribution prohibited.

01 The big picture

One pipeline, one spine, one loop. Behavioural science grounds every step; the adversary gate sits before any money moves; the Bertha codebase underneath is a **read-only reference library** — its prompts, formulas and schemas are read, never run.



02 The five minds

The skill thinks as five entangled disciplines — the same five the Bertha OS was built from. They are coupled in the references because they are coupled in reality.

STRATEGIST

Disciplined subtraction

Strategy deletes wrong options. Every deliverable shows the roads considered and rejected — that's what the client pays for.

CREATIVE

The carrier of intent

Creative is not decoration; it carries strategy into emotion. Hooks first. Anti-genericism is law: if it fits any brand, rewrite it.

MEDIA PLANNER

Physics, not channels

TikTok earns, Meta declares, Google auctions, LinkedIn buys the work-self, email is owned. Money follows proof, with kill & scale criteria.

BEHAVIOURAL SCIENTIST

Persuasion, nameable

Eight principles, each with an ethical guardrail. Two stances — guardrailed Cialdini or hard NO-FOMO — chosen explicitly, never mixed.

DATA ENGINEER

The loop is the product

RFV segmentation, churn early-warning, attribution honesty. Estimated vs measured always labelled. Marketing intelligence compounds.

03 The document hub

Documents are the working surface: parse in, judge everything, produce out. Reporting is the hub — every chart pairs with the action it implies, and every report ends with what changes next cycle.

Creative brief

in & out — job, proposition, tone matrix, acceptance criteria

Media brief

in & out — KPIs, cohorts, awareness distribution, flighting

Client documents & data

question banks · constraint detection (Critical→Low)

REPORTING the hub

campaign readout · board report
test readout · insight digest

so-what rule: chart → action

Decks

assertion headlines — the title IS the takeaway; brand-themed

PDF documents

strategy docs, audits, naming shortlists (T1–T8 skeletons)

Insight plug-ins

ad platforms · GA4 · CRM · email — one KPI/funnel contract

 Every generator has a judge — outputs are scored on five dimensions before delivery; below the bar, regenerate. Missing data is reported as missing, never interpolated.

04 Mode router

Ask in plain language; the skill routes to the engine and reads its reference before producing.

THE ASK	ENGINE	REFERENCE
Full marketing strategy / campaign plan	7-phase pipeline	<code>strategy-pipeline.md</code>
Positioning / category strategy	Dunford pipeline + judge	<code>positioning-and-offer.md</code>
Offer, pricing, tiers, discounts	Offer stack + pricing psychology	<code>positioning-and-offer.md</code>

"How much should we spend?"	Unit economics + envelope math	paid-media.md \$0
Brand growth, distinctive assets, brand health	CEPs + famexuniqueness + brand ledger	brand-growth.md
Analyze client documents & data	Extraction method	document-intelligence.md
Audit site / assets / capabilities	Front-door teardown	asset-audit.md
Cohorts, personas, audience discovery	Persona contract + discovery	cohorts-and-personas.md
Hooks, taglines, scripts, concepts	Creative engine	creative-engine.md
Names	Five-pass naming sprint	naming.md
Pressure test / red team	Adversary mode	adversary.md
Crisis response	Triage + first-hour protocol	adversary.md \$7
SEO / organic search / AI answers	Search physics	platform-physics.md \$7
Influencer / creator / UGC	Creator physics	platform-physics.md \$6
Media plan, budget, spend stewardship	Doctrine + ROI/Risk/Cost scorecard	paid-media.md
Platform specs & calendars	Platform physics + visual physics + outer ring	platform-physics.md
Persuasion scoring, brand audit	Bias radar · TMCG · Belief Stack	behavioral-science.md
CRM segments & lifecycle	RFV + churn + the loop	crm-loop.md

Email program	The owned channel	email-marketing.md
Briefs, decks, PDFs, reports	The document hub	deliverables-and-reporting.md
Deliverable skeletons	Templates T1–T8	templates.md
Deep dive into Bertha source	Read, never run	source-code-map.md

05 The operating doctrine

Eight rules, applied in every mode.

01 Strategy is disciplined subtraction

Show what was rejected, and why.

02 Anti-genericism is the law

If it could apply to any product in any market, rewrite it with specifics.

03 Every claim carries its why

Principle, data point, or platform mechanism — rationale is the deliverable.

04 Every generator has a judge

Score before delivering; regenerate below threshold. No first drafts.

05 Serve both masters

Answer the ledger and protect the magic. Flag any trade between them.

06 The loop, not the line

End by stating what measurement feeds the next brief.

07 Ethical stance is explicit

Stance A (guardrailed Cialdini) or
Stance B (NO-FOMO). Never mixed
silently.

08 Label estimates

Sizing, benchmarks, predictions —
assumptions shown, estimated vs
measured marked.

06 Using it with Claude

The skill installs once, then is driven in plain language inside Claude Code. It routes your ask to the right engine; your job is to feed it real material, give it real numbers, and let it argue. This is the difference between a demo and a machine that earns fees.



1 · Install once

clone or symlink into
`~/.claude/skills/` —
see README



2 · Just ask

any marketing question
auto-triggers it — or
invoke
**/micro-cottage-
advertising-os-
factory**



3 · Feed it

briefs · brand guidelines
· URLs ·
ad/GA4/CRM/email
exports — originals, not
summaries



4 · Let it argue

answer the intake
questions; demand the
pressure test before
money moves



5 · Close the loop

bring performance data
back — the report
rewrites the next brief

Say it in plain language — the router does the rest

SAY THIS TO CLAUDE

WHAT RUNS

WHAT YOU GET BACK

"Audit acme-widgets.com — where does the front door leak?"	Asset & website audit	Scorecard, friction & sludge ledger, 30/60/90 plan, The One Thing
"Build a full strategy from these three PDFs" (<i>attach them</i>)	Document intelligence → 7-phase pipeline	T1 strategy doc — positioning, cohorts, campaign concepts, measurement
"We have £3k/month — where should it go?"	Unit economics + media doctrine	Envelope math, derived target CPA, long/short split, T5 media plan
"Who exactly should we be targeting?"	Cohorts & personas	Cohort cards — sizing, Four Forces, hook lines, CEP tags
"Give me a campaign idea and hooks for the spring launch"	Campaign platform + creative engine	Organizing idea, scored hooks & taglines, asset matrix
"Poke holes in this plan before we spend"	Adversary mode	Red-team table, assumptions ledger, kill criteria, verdict
"Name the product"	Five-pass naming sprint	Shortlist with rejection gauntlet + clearance flags
"Write the briefs for our editor and media buyer"	Deliverables hub	T4 creative brief (+ T4-C creator variant), T5 media plan — ready to paste
"Here are last month's numbers — what now?"	Reporting hub + CRM loop	Readout with so-what actions; the re-brief for the next cycle

Getting the best out of it

FEED

Originals, not summaries

The extraction engine mines verbatim copy, tone and constraints. Paraphrases lose exactly the specifics the anti-genericism law needs.

NUMBERS

State the envelope

Budget, margin, AOV, market. Unit economics turns them into real targets — without them, everything ships labelled *estimated*.

ETHICS

Pick your stance

Stance A (guardrailed full-spectrum) or Stance B (NO-FOMO, positive-only) — say it once per engagement; it gates urgency, scarcity and loss frames.

HANDOFF

Name your platforms

Say "we produce in Canva and buy in Meta Ads Manager" and briefs arrive as ready-to-paste settings and prompts, negative prompts included.

PROCESS

Chain the modes

Audit → strategy → cohorts → creative → pressure test → media plan beats one-shot asks. Each output is the next mode's input; the brief is the root object.

GATE

Demand the adversary

Nothing that spends money or ships publicly skips the red-team panel. "Pressure test this" is the last thing you say before launch.

DIALOGUE

Push back

It shows what it rejected and why — that's the doctrine. If you disagree, argue; the rationale is part of the deliverable, not decoration.

COMPOUND

Close the loop

Bring results back in. A report that doesn't change the next brief is a receipt, not intelligence — the loop is where the machine compounds.

☞ The skill produces the **judgement** — strategy, cohorts, briefs, plans, scores, and the rationale for all of it. Production and distribution stay on your platforms; the Bertha codebase underneath is read, never run.

07 Provenance

Where the intelligence comes from — and why this repository is protected IP.

🏠 Bertha marketing OS codebase

~19 marketing microservices + frontend · ran for real clients · NOT running now — a read-only reference library of prompts, scoring formulas, cohort benchmarks, judge rubrics and platform mechanics

↓ distilled (2026) · deep dives via **source-code-map.md**

📖 This skill — seventeen reference engines

extended with Agencio ideation: naming sprints, red-team panel, email program, LinkedIn/YouTube physics, the reporting hub — then hardened by a four-seat adversarial expert review (behavioural scientist · head of media · creative director · CMO): positioning & offer engine, brand-growth layer, unit economics, incrementality, crisis protocol



📦 Client outputs, with rationale

deliverables belong to the engagement — the methodology belongs to Agencio