

The 9 Ecom Creative Skill Pack

A practical workbook for finding the real reasons people buy, building ads Meta reads as genuinely new, and deciding what to kill, scale, or hold, with Claude skills instead of guesswork.



Contents

01	Customer Language Download	2
02	Real Buyer Profile	3
03	Competitor Ad Feed	4
04	Competitor Angle Gap Pull	5
05	Hook Pattern Breakdown	6
06	Weekly Concept Engine	7
07	One Photo, Ten Ads	8
08	Asset Diet Fix	9
09	Kill / Scale / Hold Matrix	10
+	Setup SOP	11



Customer Language Download

BEGINNER: paste-and-run prompt · ADVANCED: customer-research skill, Language Download mode

WHY IT WORKS

Every hook you've ever seen that stopped someone mid-scroll was written in the customer's own words, not a marketer's. The gap between "supports healthy joint function" and "finally slept through the night without my knee waking me up" is the entire gap between an ad that gets scrolled past and one that gets clicked. Most brands never close that gap because nobody sits down and actually reads a month of reviews looking for how real customers phrase their problem.

HOW IT WORKS

This skill pulls the last 30 days of reviews, Klaviyo replies, and ad comments, and extracts the 30 phrases that show up again and again in ways that are specific to this brand's customers, not generic praise anyone could leave on any product. Each phrase gets sorted into one of three buckets: how the customer describes the problem, how they describe the result, or what almost stopped them from buying. Every phrase gets cited back to its source, and nothing gets smoothed out or made punchier along the way, the value here is the customer's exact words, not an improved version of them.

HOW TO DEPLOY IT

Run this every one to two weeks, customer language drifts as a product's audience grows and changes. Feed the output straight into hook writing and ad copy: when a hook uses an exact phrase from this download instead of a marketer's paraphrase of it, it reads as something a real person would actually say, which is most of what makes a cold audience stop and pay attention.



Real Buyer Profile

BEGINNER: paste-and-run prompt · ADVANCED: customer-research skill, Buyer Profile mode

WHY IT WORKS

A demographic persona (34-year-old woman, household income \$75K+) has never bought anything. A person with a specific trigger, a specific objection they overcame, and a specific identity they're protecting or building toward, that's who actually clicks buy. Most ad accounts are still targeting the demographic, which is why the copy feels like it's talking at a spreadsheet instead of a person.

HOW IT WORKS

This skill builds the psychographic layer sitting underneath the demographic one: what actually triggers a purchase, what objection almost stopped it, what identity claim the customer is making by buying this, what they're really aspiring toward, and just as important, who this product is clearly not for. It pulls from 90 days of real purchase, engagement, and review data, and every trait it reports has to trace back to something real in that data. A demographic fact and a psychographic one never get blended together into one vague statement, they're kept separate because they answer different questions.

HOW TO DEPLOY IT

Build this once per brand, then revisit it quarterly or whenever the customer base shifts (a new acquisition channel, a new product line, a shift in who's actually converting). This is the foundation everything else in this kit reads from, angle work, hook work, concept scoring, all of it gets sharper once there's a real buyer profile sitting behind it instead of a guess.



CHAPTER 03

Competitor Ad Feed

BEGINNER: paste-and-run prompt · ADVANCED: competitor-intelligence skill, Ad Feed mode

WHY IT WORKS

Every dollar spent testing an angle a competitor already proved, or already proved doesn't work, is a dollar that didn't need to be spent testing at all. Most brands either ignore competitors outright or scroll their feed occasionally without any structure to what they're looking at. Neither approach turns into something usable.

HOW IT WORKS

This pulls ads from tracked competitors that have been live for 30 or more days, the filter that actually matters, because an ad still running after a month has survived long enough to be a real winner, not just a test that hasn't been killed yet. Per ad, it extracts the format, the hook and what mechanism it's built on, the visual anchor, the CTA, the offer, and whatever audience signal can be inferred. Then it's cross-referenced against the brand's own live ads to see where there's overlap and where there's open space.

HOW TO DEPLOY IT

Run this weekly alongside the Weekly Concept Engine. Its output feeds directly into two other chapters in this pack, Hook Pattern Breakdown and Competitor Angle Gap Pull, both read what this skill pulls, so it should always run first. Without live Meta Ad Library access, this still works from screenshots or descriptions of ads the operator has personally noticed sticking around.



Competitor Angle Gap Pull

BEGINNER: paste-and-run prompt · ADVANCED: competitor-intelligence skill, Angle Gap mode

WHY IT WORKS


Two kinds of mistakes happen with angles: testing one a competitor's already proven works (expensive to discover independently), and testing one nobody's touched because it doesn't actually work (expensive to discover the hard way). The useful middle ground, an angle the data supports but nobody in the category has claimed, is easy to miss without actually mapping the category first.

HOW IT WORKS

This catalogs every competitor angle by which specific customer tension it's pulling on, then checks that catalog against the brand's own live angles. What's left over after that comparison is the real gap list, angles the ICP data supports that nobody nearby is running. Each gap gets three concept variants built around it, and just as importantly, gaps that are technically open but would require a claim the product can't actually back up get flagged separately and left alone.

HOW TO DEPLOY IT

Run right after Competitor Ad Feed, in the same session, it reads that skill's output directly. Cross-reference the gaps this produces against the Real Buyer Profile, a gap that doesn't connect to a real ICP tension isn't worth chasing even if the data shows nobody's claimed it yet.



Hook Pattern Breakdown

BEGINNER: paste-and-run prompt · ADVANCED: competitor-intelligence skill, Hook Pattern mode

WHY IT WORKS

There are roughly ten recurring hook mechanisms (contrarian, curiosity, number-led, named-enemy, physical demo, confession, reframe, stakes, reversal, ridiculous), and most categories end up over-relying on two or three of them without anyone noticing, because nobody's actually counting. The mechanisms nobody's using aren't necessarily bad, they're just unclaimed.

HOW IT WORKS

This groups every competitor hook pulled by the Ad Feed into its underlying mechanism, then flags which ones are over-indexed across the category. It cross-references against the brand's own hook library if one exists, and scores the findings by two separate things: how saturated a mechanism already is, and how novel it would be if the brand ran it. A mechanism that's saturated but still clearly working is a different finding than a mechanism nobody's tried, both matter, but for different reasons.

HOW TO DEPLOY IT

Run this after Competitor Ad Feed, same session. The mechanism gaps this surfaces are some of the strongest raw material to hand to the Weekly Concept Engine, a hook mechanism the whole category is ignoring is often the fastest way to make a new concept feel genuinely different in a crowded feed.



Weekly Concept Engine

BEGINNER: paste-and-run prompt · ADVANCED: concept-engine skill, Weekly Concept Engine mode

WHY IT WORKS

The single most common reason testing volume stays low isn't budget, it's that generating twenty genuinely distinct concepts by hand, every single week, is slow and most people run out of real ideas after four or five. The algorithm needs new creative constantly. A blank page every Monday is how testing volume quietly dies.

HOW IT WORKS

This generates 20 distinct concepts for one hero product in a single pass. Every concept gets a premise specific enough to storyboard directly, a hook under 12 words with the pattern break front-loaded, a format, a funnel position, the specific ICP tension it targets, and two scores, fit against the ICP and novelty against what's already been tested. The 20 get ranked by combined score, and the top 10 become the week's actual testing batch. Nothing gets handed over as a near-duplicate of something else on the list, every concept has to be structurally distinct from the other 19.

HOW TO DEPLOY IT

Run this weekly, it's built for a recurring cadence, not a one-time brainstorm. Feed it everything the earlier chapters in this pack produce, customer language, buyer psychology, competitor gaps, hook mechanisms, the richer that input, the sharper the 20 concepts come out. The Kill/Scale/Hold Matrix (chapter 9) is what should eventually tell you whether last week's picks from this engine were the right ones.



One Photo, Ten Ads

BEGINNER: paste-and-run prompt · ADVANCED: one-photo-ten-ads skill (standalone)

WHY IT WORKS

A traditional shoot for ten ad variations is a two-day production with ten outfit or setup changes, running \$5,000 to \$8,000. Most brands can't test video creative at the volume the algorithm actually rewards because producing it the traditional way is simply too slow and too expensive to do weekly.

HOW IT WORKS

This is a two-step pipeline. Step one takes a single product photo and generates ten genuinely distinct "ways" to feature it, specific use cases, styling angles, occasions, or applications, following the viral "one product, ten ways" reel format. Step two takes those ten variations plus the product photo (and a character reference, if one's being used) and assembles one complete, paste-ready video generation prompt: an opening hook, a main montage running through all ten variations, on-screen text per variation, consistent camera and lighting direction, and a closing reveal with a call to action.

HOW TO DEPLOY IT

About 15 minutes of operator time and \$12 to \$20 in render credits for the finished output, instead of a multi-day shoot. The finished prompt is paste-ready for Seedance 2.5 (through Higgsfield or direct access), and the rendered result still needs a pass through CapCut for music, text overlay, and final export. This is one of the fastest ways in the entire kit to turn a single piece of existing creative asset (one photo) into a week's worth of testable video.



Asset Diet Fix

BEGINNER: paste-and-run prompt · ADVANCED: asset-diet-fix skill (standalone)

WHY IT WORKS


Meta's Andromeda algorithm groups ad variations it judges structurally similar into a single "entity ID." This means ten different hooks on the same underlying concept can still get treated as one ad competing for one audience pool, reach capped, no matter how different the words feel to a human reading them. This is specifically why a winning ad's CPMs start climbing even while new hook variations keep getting added.

HOW IT WORKS

The fix isn't more hooks, it's making a smaller number of hooks structurally distinct enough that Andromeda reads each one as separate inventory. This chapter's skill takes a brand's proven ICP tensions and builds a relevance map of pains, triggers, and aspirations, narrows that down through a filtering sequence to the four strongest hook candidates, then wraps each of those four in one of four structurally different formats: a seated two-person podcast conversation, a street interview, a solo walk-and-talk, or a first-person POV. Same underlying concept, visually and structurally distinct execution.

HOW TO DEPLOY IT

Use this on proven winners that have hit a reach ceiling, not on brand-new concepts, it's a reach-multiplication tool, not a concept-generation one. In a tested sample of 30 DTC brands, ten hook variations alone got grouped into one entity ID (one audience pool, 28 to 32 percent average hook rate, 1.4x average ROAS); the four structural wraps reached four distinct audience pools (45 percent or higher average hook rate, 3.5x to 6x average ROAS, 38 percent lower average CPA). The four finished briefs this produces still need a downstream production tool, MakeUGC, Maxfusion, or Higgsfield, to actually film or render them.



Kill / Scale / Hold Matrix

BEGINNER: paste-and-run prompt · ADVANCED: concept-engine skill, Kill/Scale/Hold mode

WHY IT WORKS

Most of the money lost in a testing budget doesn't come from bad concepts, it comes from good decisions made too late: a winner that should have been scaled three days ago, or a loser that kept absorbing spend for a week past when the data already made the call obvious. Guessing on budget moves, even educated guessing, is slower and more expensive than applying a consistent rule.

HOW IT WORKS

This pulls the last 24 to 48 hours of data per live ad and sorts every one of them into exactly one of four categories: kill, scale, hold, or iterate. For anything marked scale, it calculates a specific budget increase based on the ad's maturity and scaling history, smaller increases for a first scale on a young ad, larger ones for a proven ad that's already absorbed increases before. For anything marked kill, it logs a likely root cause so the pattern can inform what gets tested next. Nothing gets auto-executed, every recommendation goes to the operator for approval before any budget actually moves.

HOW TO DEPLOY IT

Run this daily once ads are live, it needs at least 24 hours of data per ad to make a reliable call. If there's no defined kill, scale, or hold threshold yet, this chapter's skill will build a starting set from a target CPA and target ROAS rather than stalling, those thresholds get saved so every future run has them on file. This is the chapter that closes the loop on chapter 6, it's how you find out whether last week's concept picks were actually the right ones.

Setup SOP

BEGINNER TIER (INCLUDED WITH EVERY KIT)

Create a Project in Claude, paste the Router system prompt into Custom Instructions, upload the Knowledge Base file as Project knowledge, then just talk to it in plain language inside that Project. No skill names to memorize, no prompts to copy by hand, the router figures out which of the nine does what you're asking for. Full step-by-step is in the Setup Guide that shipped with your kit.

ADVANCED TIER (CLAUDE SKILLS PACK)

The same nine deliverables from this workbook, rebuilt as five real Claude Skills (bundled by workflow rather than one-to-one, since several of these are naturally run together). Download the Advanced Skill Pack, unzip it, and drop each .skill file into Claude Code's or the Claude Desktop app's skills directory. Once installed, Claude reaches for the right skill automatically based on what you ask for, you don't open a Project or paste a router prompt at all, you just work.

Either tier runs the same underlying system. The beginner tier is the fastest way to get started with zero setup beyond one copy and paste. The advanced tier is for anyone already comfortable in Claude Code or the Desktop app who wants these to trigger automatically alongside everything else they're doing, instead of living inside one dedicated Project.