

Don't trust the *output*. Trust the receipts.

Brand: _____
Date: _____

AI will generate personas confidently whether the data supports them or not. **Your job is to demand proof.** The five checks are here. The prompts to run them are on page 2.

HOW WE USED TO DO IT VS. WHAT AI CHANGES

→ AI compresses the synthesis step. It doesn't replace the research.

THEN: MANUAL RESEARCH STACK (IN ORDER)	NOW: AI DOES THE SYNTHESIS
<ol style="list-style-type: none"> Reddit first. Unfiltered. Nobody's selling you anything. Search the category, read the threads, screenshot the exact language people use. Reviews second. Purchase-validated voice of customer. What do people say after they bought? What triggered them? Competitor ads last. By now you know what your real buyer sounds like. You can spot whether a competitor is actually speaking to them. <p><i>The order matters. Each layer informs how you read the next one.</i></p>	<ol style="list-style-type: none"> Do the same research. Reddit, reviews, competitor ads. AI doesn't do this part — it analyzes what you give it. Paste the raw material in. Don't summarize it first. Give Claude the actual text so it finds patterns in real language. AI finds themes faster. What used to take hours now takes minutes. But output is only as good as what you fed it. <p><i>Garbage in, garbage out. Still applies.</i></p>

5 WAYS TO CHECK THE OUTPUT

Run these after AI gives you personas. Don't move to pillars until all five pass.

<p>01</p> <p>ASK FOR RECEIPTS</p> <p>Demand real quotes from the source material for every persona before you trust it.</p> <hr/> <p>If it can't produce them, the persona isn't grounded. If the quotes don't match what Claude claimed, same problem.</p>	<p>02</p> <p>CHECK SAMPLE SIZE</p> <p>AI builds confident personas from 4 reviews or a \$200 ad set. That's a coincidence, not a pattern.</p> <hr/> <p>Reviews: fewer than 20–30 is too thin. Spend: under ~\$500–1k has no statistical weight. Don't build a pillar on it.</p>	<p>03</p> <p>CHECK A SECOND AI</p> <p>Paste the same raw material into ChatGPT or Gemini. Ask for the same output. Compare.</p> <hr/> <p>Both agree = real signal. They disagree = your data is ambiguous. Look at it yourself before you build on it.</p>	<p>04</p> <p>CHECK FOR A TRIGGER</p> <p>A real persona has a trigger — the moment that pushed someone to start looking for a solution.</p> <hr/> <p>"Cares about health" = demographic. "Just found out her kid has eczema and CVS isn't cutting it" = trigger. That's an ad.</p>	<p>05</p> <p>CHECK FOR SKU BLENDING</p> <p>Multi-product brands: AI averages buyer profiles into one persona that doesn't describe anyone accurately.</p> <hr/> <p>Protein powder buyer and pre-workout buyer are two different people. A blended persona briefs neither correctly.</p>
---	---	--	---	---

The rule: AI sounds confident whether the data backs it up or not. Real signal produces real quotes. If Claude has to reach, you'll see it in the reaching.

The *prompts*. Copy, paste, run.

Run these in order after AI generates your personas. Don't skip ahead.

Each prompt maps to one of the five checks on page 1. Paste your source material into Claude first, get the personas, then run these. **Don't accept a persona until it survives all five.**

01 ASK FOR RECEIPTS
Run this immediately after getting personas

PROMPT

"For each persona you identified, give me 3 direct quotes from the source material I provided that support it. Quote the actual text — do not paraphrase. If you cannot find 3 real quotes for a persona, tell me that instead of generating examples."

If Claude generates quotes that don't exist in your pasted material, the persona was invented. Reject it and ask it to try again using only what's in the source.

02 CHECK SAMPLE SIZE
Verify data volume before trusting any output

PROMPT

"For each persona you identified, tell me how many individual reviews or data points from my source material support it. Flag any persona supported by fewer than 10 distinct sources and label it low-confidence."

Under 20–30 reviews = too thin to call dominant. Under ~\$500–1k ad spend = no statistical weight. Note it — don't build a pillar on a hunch.

03 CHECK A SECOND AI
Paste this into ChatGPT or Gemini with the same source

PROMPT

"Here is raw customer research for [BRAND]: [paste same source material]. Identify the 2–3 most distinct customer personas. For each, describe: the trigger that brought them to this product, the language they use to describe their problem, and the outcome they want. Do not use demographic labels as persona names."

Compare side by side. Both surface the same trigger = real signal. Completely different personas = ambiguous data. Go look at it yourself.

04 CHECK FOR A TRIGGER
Push back if AI gave you demographics instead

PROMPT

"Review the personas you gave me. For each one, identify the specific trigger moment — the event or life change that pushed this person to start looking for a solution. If a persona only has demographic descriptors without a clear trigger, rewrite it to be trigger-based using evidence from the source material."

If Claude can't find a real trigger in the data, the persona isn't ready to brief against. Don't move forward until it has one.

05 CHECK FOR SKU BLENDING
For any brand with more than one product — run this every time

PROMPT

"Review the personas you identified. For each persona, specify which product or SKU the supporting quotes came from. If a persona draws from reviews or data across multiple products, flag it as a blended persona and rebuild it separately for each product."

WHY THIS MATTERS

The protein powder buyer and the pre-workout buyer have two different triggers and two different buying contexts. AI will blend them into one averaged persona that doesn't accurately describe either — and definitely won't produce a useful brief. If Claude can't tell you which SKU each supporting quote came from, the persona needs to be rebuilt per product before you use it for anything.