

**HOW TO  
BUILD  
CLAUDE  
SKILLS  
THAT  
ACTUALLY  
*WORKS.***

FREE CHAPTER — A JOURNEY WITH AI PUBLICATION

# **How to Build Claude Skills That Actually Works**

*Chapter 1 • The Skill That Did Nothing*

[journeywithai.com](https://journeywithai.com)

# 1

## CHAPTER 1

# The Skill That Did Nothing

---

**C**lass, let us begin.

There is a moment I want you to picture.

You are sitting in front of your laptop. Late evening. The evening where the room is quiet and the only sound is the small electric hum of a fan or a fridge somewhere behind you. You have just spent two hours writing what you are calling your "brand voice skill." A long markdown file. Headings. Bullet points. A list of words you would never use. A list of words you love. You have written

the tone profile. You have written the formula. You have written, carefully, "the brand should feel warm but not soft, confident but not arrogant, witty but not flippant."

You have saved it. You have installed it. You have opened a fresh Claude window and typed the prompt you would normally hand to your in-house copywriter.

"Write me a LinkedIn post about onboarding for our SaaS platform."

You hit enter.

Three seconds later, Claude produces a perfectly polished, perfectly readable, perfectly forgettable post. Generic hook. Generic value statement. Generic call to action. Three emojis. The word *empower* somewhere in the second paragraph.

You read it once. You read it again. You sit with it.

Then you do the thing every founder does at this exact moment. You open the skill file you just spent two hours on. You scroll through it. You add another line. You add a sterner instruction. "ALWAYS use specific brand examples." "NEVER use the word empower." You save. You close the file. You re-run the prompt.

The new post is ten percent better.

It is still not yours.



# The lamp that broke me

I have watched dozens of founders and operators sit in this exact chair. I have sat in it myself. The first time it happened to me, I had just finished writing a voice skill for **Eyda Homes**, the 25-year-old artisanal home brand my partner and I run as co-founders out of India. The skill was thirty pages long. I had poured everything into it. The colour palette. The tone notes. The vocabulary preferences. The audience definitions for India versus the UK. I knew the brand inside out. I had been on the factory floor, watched the carpenters work, listened to the founders' father talk about how they sourced wood in the seventies. Nobody knew this brand the way I did.

Then I asked Claude to write a product description for a new lamp.

It produced two paragraphs about "warm ambient lighting" and "elevating your home" and how "this beautifully crafted piece will transform your space."

I closed my laptop and walked out of the room.

The post that day, the lamp description that night, the LinkedIn caption a friend asked me to fix the next morning. They were not the same content. They were the same problem.

The skill had words. It did not have architecture.



# The pattern I had already seen

Here is the strange part. I should have recognised the pattern instantly. I had seen it for ten years before I ever opened a SKILL.md file.

I built **Blushush** in London in 2022. Digital marketing agency. We worked with clients across seven countries. **N1 Payments** was with us for seven years. **ARCC Bikes** for four. **Gunpowder Restaurants** built their digital across seven locations between London, Lisbon, and India with us. **Born Clothing** ran their full revamp through us. The agency had range.

What we did not have was a way to make clients arrive prepared.

A new client would book a call. We would pitch. They would say yes. They would send over what they were calling a brief. Most of the time, the brief was three paragraphs of taste statements and one sentence that said something like, "We want it to feel premium but also approachable, modern but not too modern." A list of brands they admired, none of which actually resembled the brand they themselves were building. A few colour preferences. The word *innovative*, somewhere.

We would design something. Send it across. The client would nod. Sign off. Pay the invoice.

Six months later, I would hear the same line for the dozenth time. *We are looking for a new agency. Something feels off. We cannot name it.*

I knew exactly what was off. The work was based on what the client had told us, which was a collection of preferences that were not connected to any underlying thinking. We had no choice but to fill in the gaps. So we did. With our best guesses. Sometimes the guesses were close. The guesses were never the brand, because the brand had not yet been defined, and you cannot guess a thing that does not yet exist.

The disconnect was never between the agency and the client. It was between the client and themselves.

By 2024, when I started writing my first Claude skills for these same clients, I was still operating with the agency-side instinct. I was writing skills the same way I had been writing creative briefs for a decade. *Tone: warm but confident. Audience: founders. Colours: navy, gold. Avoid: corporate jargon.*

The skills produced exactly what those briefs had produced for ten years. Output that was technically on-brand, in that nothing in it directly contradicted the document. Output that was actually generic, in that nothing in it was specifically the brand either.

I had imported the same problem from the agency world into the AI world, and called it a different name.



## **The realisation, named plainly**

Here is the part nobody talks about, because nobody wants to.

When you write a Claude skill the way you would write a creative brief, you are doing exactly what an unprepared client does when they walk into an agency for the first time. You are bringing preferences. You are bringing adjectives. You are bringing a list of things you like and a list of things you do not. What you are not bringing is the thing the work actually needs.

You are not bringing the *thinking behind the preference*.

This is the entire problem in one sentence. And once you see it, you cannot unsee it.

You wrote *the brand should feel warm but not soft*. What does *warm* mean for this specific brand? Where does the warmth come from? Is it the warmth of a craftsman who knows your name, or the warmth of a friend who knows what you have just been through, or the warmth of a teacher who is glad to see you in class today? All three are warm. All three sound completely different on the page. If you do not know which one your brand is, Claude will produce the average of all three. The average of three different warmths is a generic warmth. Generic warmths sound like AI.

You wrote *the brand is confident but not arrogant*. Where does the confidence come from? From decades of work? From a lived experience that nobody else in the field has? From a refusal to apologise for an inconvenient truth? Each of these produces a different confidence. Without the source, Claude reaches into the median of all confidences it has ever read on the internet, and produces the average. *We empower brands to unlock their full potential*. That is the average. That is what the median sounds like.

A list of preferences gives Claude a list of preferences.

The architecture is what produces the brand.



## What architecture actually means

The word *architecture* does a lot of work in this book. Let me unpack it once, here, so the rest of the book has a foundation to build on.

When I say a skill needs architecture, I mean five things, and most skills are missing four of them.

### **One. The skill knows what skill it is.**

A voice skill and an output skill are not the same thing. A voice skill teaches Claude how to sound. An output skill teaches Claude how to produce a specific deliverable, every time. A framework skill teaches Claude how to think through a class of problem. An identity skill encodes the soul beneath everything. These are different jobs. They take different ingredients. Mixing them is the most common skill-writing mistake. The skill that did nothing on my Eyda lamp was a voice skill that was trying to be an identity skill, an output skill, and a framework skill all at once. It was not good at any of them.

### **Two. The skill has a description that actually triggers.**

The description field at the top of a SKILL.md file is the single most undervalued part of any skill. Most people write it like

marketing copy. They are writing it for the wrong audience. The description is read by the routing system that decides whether to load the skill in the first place. A perfect skill that does not load is invisible. A mediocre skill that loads reliably will outperform it every time. Most failures I see in real skill libraries trace back here. The skill exists. The skill is good. The system never loaded it.

### **Three. Every rule has a reason.**

A skill that says *ALWAYS use specific brand examples* is brittle. The next edge case it does not anticipate, it will break. A skill that says *Specific brand examples make the writer sound like they actually live inside the brand. Generic founder stories make the writer sound like they have read about the brand. Use the specific example whenever a real one is available, and tell me when one is not* is robust. Reasoning beats rules. The model is smart. Treat it that way.

### **Four. The voice is anchored in real moments.**

Every tonal attribute the skill names should trace back to a lived experience. If a brand is sarcastic, the skill should explain *why*. The sarcasm is not a stylistic choice; it is the residue of having watched the same preventable mistake happen a hundred times. If a brand is warm, the skill should explain where the warmth comes from. The warmth is not decoration; it is the founder remembering what it felt like to be alone with a problem nobody would name for them. This is the section nobody else writes. This is what makes the difference between hollow imitation and recognisable voice.

### **Five. The skill knows what to refuse.**

Specific patterns the brand will not produce. Specific words it will not use. Specific structures that mean the draft has failed. The brand is partly defined by what it is. The brand is also defined by what it refuses to be. A skill without a refusal list will eventually drift into producing the average of everything it has ever read. The refusal list is the immune system.

If a skill has all five. The right type, a description that triggers, reasoning behind every rule, voice anchored in real moments, and a clear list of what it refuses. Claude can produce work that sounds like the brand. Not always. Not perfectly. But reliably enough that you will start sending the work without rewriting it.

If the skill is missing two or more of these, the output will drift toward generic, and you will find yourself, late on a Tuesday, adding another rule to a file that already has too many rules, and watching the output get ten percent better and still not be yours.



## **The single line that anchors this book**

So what is the architecture?

The whole rest of this book is the answer. But before we get to any of it, I want to give you the single most important shift in thinking you will make as a skill writer. If you take only one thing from this chapter, take this.

A skill is not a list of rules. A skill is a frame.

A list of rules tells Claude what to do. A frame teaches Claude how to think.

Rules are what you write when you do not yet know your brand. Frames are what you write when you do. Rules say *always use the word clarity*. Frames say *this brand exists because its founder watched 150 businesses open and close in his childhood home, and the one variable that decided which ones survived was clarity. The word clarity is not a stylistic choice. It is the brand's diagnosis of every problem in business. Use it whenever a problem is being diagnosed. Avoid it when you are stating a feature.*

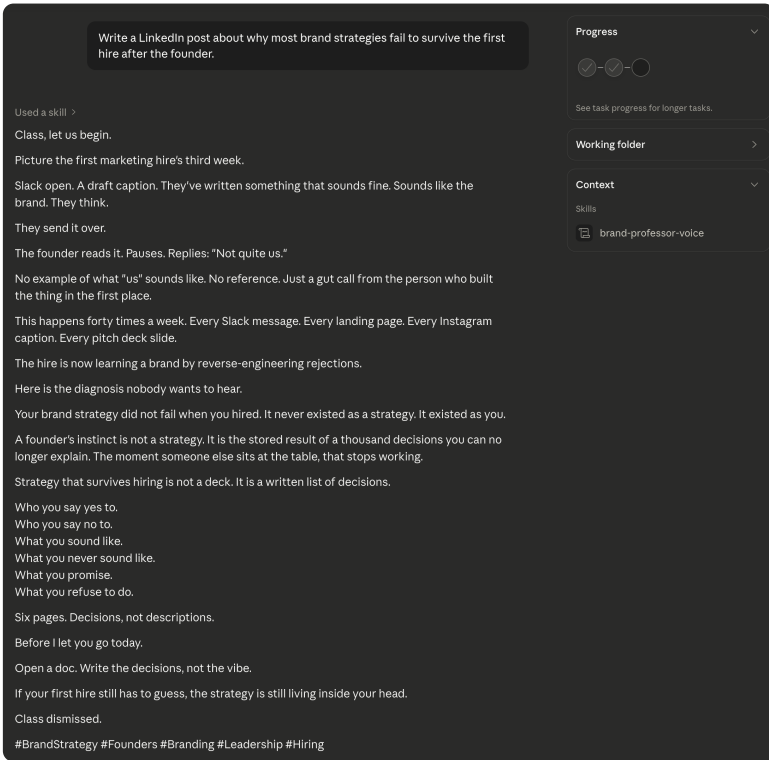
One of those is two clauses long. The other is two paragraphs. The two paragraphs is the one that actually changes the output.

Rules pile up. Frames teach. The skills that work are the ones where Claude can answer the prompt from inside your head, not from inside its training data.

Read that line again.



*A skill works when Claude can answer it from inside your head, not from inside its training data.*



**Figure 3. One skill, one prompt. The Brand Professor voice skill firing on a LinkedIn request. The skill loaded indicator (top right) shows brand-professor-voice as the only skill active. The post that follows reads like the author wrote it himself, late on a Tuesday.**

That is the line. Everything in this book is about how to write a skill that puts Claude inside your head. Which means everything in this book is about doing the thinking first.



## **A quick glimpse of what the rest of the book does**

Before I let you go today, let me show you the road.

Part 2 of this book introduces the Brand Professor Skill Operating System. Six types of skills. Three categories. The taxonomy nobody else has named. By the end of Part 2, you will be able to look at any operational frame you might want to encode and identify which type it is, what ingredients it needs, and which other skills it should chain with.

Part 4 is the proof. Four case studies from my own library. Brand Professor itself, the three-layer stack of identity, voice, and memory. The Blushush proposal skill, which turns a fifteen-minute conversation into a branded landscape PDF that lands in clients' inboxes. The Strategic Thinking Master Operator, the orchestrator that routes between every other strategic skill in the system. And three voice skills built on the same scaffolding for three radically different brands. Eyda Homes, Empyrean Infotech, and Ohh My Brand. You will see exactly how the framework is applied in practice, on real businesses, with real outcomes.

Part 5 is the ship-it part. How to test a skill. How to fix it when the description does not trigger. How to audit it monthly so it stays sharp.

Part 6 is the giveaway. The companion Skill Builder skill. You install it. It runs the whole process for you. You answer the questions. It builds your skill. You walk away with a working .skill file you can install anywhere.

That is the road. We are at the start of it. The chapters from here get more concrete, more specific, more usable. By the time you are halfway through Part 3, you will be writing your first skill alongside the book.



## **What to do before Chapter 2**

Before I let you go today.

Here is the small thing I want you to do.

Open the skill you have already written. Or, if you have not written one yet, open a blank file and write the skill you have been planning to write. Three pages, five pages, one page. Whatever you have.

Read it back as if you were Claude.

Now ask yourself one question. Just one.



*Could a stranger, given only this document, do the work the way I would do it?*

---

Not the way someone competent would do it. Not the way an agency would do it. Not the way "a writer with this brand voice" would do it. The way *you* would do it. The way that is recognisably you in a way nobody else could fake.

If the answer is yes, you are ahead of most of the world. The rest of this book will sharpen what you have.

If the answer is no, do not change a word yet. Keep the skill as it is. We are about to find the gap between what you have written and what the work actually requires. Once we find the gap, we will close it. The gap has parts. The parts have names. The next two chapters give them to you.

Class dismissed.

END OF FREE CHAPTER

**Twenty-one chapters. 374  
pages. The full  
methodology.**

Read the rest at the full price the book holds  
because most readers finish it and start  
writing skills the same week.

**\$29 · digital edition**

[journeywithai.com/the-book](https://journeywithai.com/the-book)

A Journey with AI publication · 2026