

FREELANCING & BUSINESS SKILLS

Selling AI Prompts

How to build a small, low-cost income stream selling AI prompts
- from testing a prompt properly through packaging, pricing,
choosing a marketplace, and a 30-day plan to your first sale.

10 Chapters

Philip Osztromok

Generated with Claude

CHAPTER 1: THE OPPORTUNITY

Selling AI Prompts

Course 1 · Chapter 1 · The Opportunity

A prompt marketplace sells something oddly abstract: not an image, not an essay, not code — just the instructions that reliably produce one. That's a genuinely strange product at first glance, since anyone can type into ChatGPT or Midjourney for free. This chapter looks at why people pay for it anyway, what a realistic income from it actually looks like, and who's actually buying.

\$0

START-UP COST

Hours

TO FIRST LISTING

~20%

TYPICAL MARKETPLACE FEE

\$1.99–\$9.99

TYPICAL PRICE PER PROMPT

What You're Actually Selling

Nobody is buying words. They're buying a **tested, reliable result** — the version of the prompt that already survived the trial-and-error most people give up on halfway through. A single line typed into an image generator is a starting point, not a finished product; a genuinely sellable prompt has already been through dozens of iterations to nail down exact phrasing, parameter values, negative prompts, and edge cases that break it.

What a Buyer Could Write Themselves

- A rough, one-line description of what they want
- No idea which keywords actually move the output
- Inconsistent results run to run

What a Tested Prompt Actually Delivers

- Wording refined across dozens of test runs
- Placeholders/variables for the buyer's own specifics
- Consistent output quality run to run

- No sense of which settings/parameters matter
- Time cost: hours of trial and error to get close

- Documented settings, model version, and known edge cases
- Time cost to the buyer: minutes

What the Market Actually Looks Like

Prompt marketplaces split into two genuinely different products. **Image-generation prompts** (for tools like Midjourney or Stable Diffusion) dominate marketplace volume, for a simple reason: a buyer can see exactly what they're getting from a preview image before paying — the product sells itself visually. **Text/workflow prompts** (for ChatGPT, Claude, and similar) sell less on visual proof and more on a documented, repeatable process — a multi-step prompt chain for something like structured customer-support replies, a resume rewrite template, or a code-review checklist. Chapter 2 breaks both categories down properly.

SET EXPECTATIONS HONESTLY

This is a genuinely low-cost, low-effort way to start — but it's also a lower-ceiling opportunity than a service business like freelance writing. Individual prompts typically sell for a few dollars each, so meaningful income depends on volume (a catalog of many prompts) rather than a handful of high-value sales. It's best framed as a passive-ish side income stream, not a fast route to a full-time replacement income on its own.

Who Actually Buys Prompts?

Marketers & social media managers

Want a consistent visual style across posts without hiring a designer.

Small business owners

Want a working ChatGPT workflow (e.g. customer replies) without learning prompt engineering themselves.

Other content creators

Want reliable image styles or scriptwriting templates to speed up their own production.

Hobbyists & students

Want interesting, well-crafted image prompts for personal projects without the trial-and-error time.

A note on platform-specific figures

Marketplace fees, typical price points, and which platforms are currently active all shift over time. Treat the numbers in this chapter (and the rest of this course) as a realistic starting reference, and re-verify current fees and listings directly on whichever platform you actually use before pricing anything.

Coming up next

Chapter 2 breaks down the three main prompt categories in detail — text, image, and workflow prompts — and where genuine buyer demand actually concentrates within each.

CHAPTER 2: TEXT, IMAGE & WORKFLOW PROMPTS

Selling AI Prompts

Course 1 · Chapter 2 · Text Prompts, Image Prompts & Workflow Prompts

Chapter 1 treated "a prompt" as one thing. It isn't. Three genuinely different products live under that label, each with a different buyer, a different way of proving quality before the sale, and a different realistic price. Picking which one to focus on shapes almost everything in the rest of this course.

The Three Categories



Image-Generation Prompts

HIGHEST MARKETPLACE VOLUME

Prompts for tools like Midjourney, Stable Diffusion, DALL-E, or Leonardo AI — a precisely tuned description plus style keywords, aspect ratio, weights, and negative prompts that reliably produce a specific visual look. This category dominates marketplace listings for a simple reason: a buyer can see exactly what they're getting from a preview image before paying. The product markets itself.

Typical buyer: content creators, marketers, hobbyists wanting a consistent aesthetic

Proof of quality: the preview image itself — no explanation needed



Text / Chat Prompts

HARDER TO PROVE, REAL DEMAND

A single, well-engineered prompt for a specific text task — a cold-email generator, a product description generator, a study-notes summariser — built for ChatGPT, Claude, or Gemini. There's no single "picture" to sell here, so quality has to be demonstrated with worked before/after examples instead of a preview image. That extra friction is exactly why this category has less competition than image prompts, even though genuine demand for it is strong.

Typical buyer: professionals and small business owners solving one repeatable task

Proof of quality: a documented before/after example, not a single glance

Workflow / Chained Prompts

HIGHEST PRICE CEILING

A multi-step sequence of prompts designed to be run in order — a five-step content-repurposing workflow, a customer-support ticket triage sequence — often packaged as a custom GPT, a Claude Project, or a Notion AI template rather than raw text. This solves a whole process, not one output, which is why it commands the highest per-sale price of the three — but it also needs the most sophisticated packaging: step-by-step instructions, not just a single prompt to paste in.

Typical buyer: small teams wanting an entire process solved, not just one task

Proof of quality: a worked example of the full process, start to finish

How They Actually Compare

CATEGORY	EASE OF PROVING QUALITY	TYPICAL PRICE	COMPETITION
Image-generation	Very easy (preview image)	\$1.99–\$4.99	High
Text / chat	Moderate (worked example)	\$3–\$9.99	Medium
Workflow / chained	Harder (full walkthrough needed)	\$9–\$40+	Lower

WHERE GENUINE DEMAND ACTUALLY CONCENTRATES

High competition in image prompts doesn't mean it's a bad category — it means proof of quality is cheap to produce, so more sellers can compete. Text and workflow prompts have a real, underserved gap precisely because proving quality takes more effort. If you're willing to build a genuine worked example, that's the lower-competition, higher-price-ceiling lane.

None of this means you have to pick only one category forever — many sellers eventually build a catalog spanning two or three. But starting with a single, deliberately chosen category (Chapter 4 covers how to pick one) keeps the early workflow manageable.

Coming up next

Chapter 3 gets specific about what separates a prompt that actually sells from one that doesn't — regardless of which of these three categories it's in.

CHAPTER 3: WHAT MAKES A PROMPT ACTUALLY SELLABLE

Selling AI Prompts

Course 1 · Chapter 3 · What Makes a Prompt Actually Sellable

A "lucky" prompt — one that happened to produce a great result once — is not a product. A sellable prompt has to survive being run by a stranger, with different inputs, on a different day, and still work. This chapter covers the four things that separate the two.

The Four Pillars of a Sellable Prompt

1. Reliability across runs

Tested across many runs (and, for image prompts, many seeds) — not judged on the one great output that happened to come out first.

2. Reusable placeholders

Built with clearly marked variables (a name, a topic, a tone) the buyer swaps in for their own use case, not baked-in specifics only useful to you.

3. Guardrails

Explicit constraints that stop the output going off the rails — a word count, a format, things to avoid, things never to invent.

4. Documentation

Which exact model/version it was tested against, and any known edge cases or settings that matter.

A Worked Example: Naive vs. Sellable

Same underlying task — a product description generator — at two very different levels of readiness.

NAIVE (NOT SELLABLE)

Write a product description for a pair of headphones.

This "works" in the sense that it produces text. But run it five times and you'll get five different lengths, tones, and levels of accuracy — and it only describes headphones. A buyer can't reuse it for anything else without rewriting it from scratch, which defeats the entire point of buying it.

ENGINEERED (SELLABLE)

You are an expert e-commerce copywriter. Write a product description for
[PRODUCT_NAME],
a [PRODUCT_CATEGORY].

Target audience: [AUDIENCE]

Tone: [conversational / luxury / technical]

Key features to highlight:

- [FEATURE 1]
- [FEATURE 2]
- [FEATURE 3]

Rules:

- 100–150 words
- No generic filler ("this amazing product will change your life")
- Include one specific, concrete sensory detail
- End with a single clear call to action
- Do not invent specifications not provided above // the guardrail against hallucinated features

The second version works for any product, in any category, for any client — that's what turns one good result into a repeatable product someone will pay for more than once.

Self-Check Before Listing

- Have I run this at least 5–10 times with different inputs, not just once?
- Does every variable a buyer would realistically want to change have a clearly marked placeholder?
- Have I told it explicitly what *not* to do, not just what to do?
- Have I written down which model/version I tested it against?

- If I handed this to a stranger with zero context, could they use it correctly on the first try?

A one-off screenshot is not proof

A single impressive output screenshotted from your best run tells a buyer almost nothing about whether the prompt is actually reliable. Chapter 5 covers a proper testing workflow for confirming consistency before you ever list anything.

Coming up next

Chapter 4 covers picking an actual niche to focus on — how to spot genuine demand and avoid an already oversaturated category.

CHAPTER 4: FINDING A PROFITABLE NICHE

Selling AI Prompts

Course 1 · Chapter 4 · Finding a Profitable Niche

"I'm going to sell prompts" isn't a niche — it's a category with thousands of competitors. The sellers who actually make consistent money picked something specific enough that they could become genuinely good at it, and specific enough that a buyer searching for exactly that thing finds them instead of getting lost in a sea of generic listings.

Reading Demand Signals

Marketplace bestseller lists

What's already selling repeatedly, not just what's listed — a listing with many reviews is a real demand signal.

Gaps in existing listings

Search a specific use case and see how thin the results are — a popular general category with almost nothing narrow is a genuine opening.

Community requests

Recurring questions in AI-focused subreddits, Discord servers, and forums — "does anyone have a prompt that actually does X" is a direct signal.

Your own professional background

A niche tied to work you've actually done (real estate, HR, e-commerce, teaching) gives you a credibility edge no amount of research alone provides.

Broad vs. Right-Sized Niches

Too Broad (oversaturated)

"Fantasy portrait prompts," "write me a blog post," "LinkedIn post generator."
Enormous existing supply, near-zero

Right-Sized (findable, provable)

"Property listing description prompts for UK estate agents," "LinkedIn thought-leadership prompts for B2B SaaS founders," "children's-book-style

differentiation, and pricing gets driven down by sheer competitor volume.

illustration prompts for self-publishers." Specific enough that a search for it actually surfaces you, and specific enough that you can demonstrate real expertise rather than generic competence.

YOU DON'T NEED THE BIGGEST NICHE — YOU NEED A PROVABLE ONE

A smaller, specific niche where you can produce genuinely convincing worked examples (Chapter 3) will consistently outsell a huge, generic category where a buyer has no way to tell you apart from a thousand other listings. Depth beats breadth here.

A Simple Niche-Selection Checklist

- Can I name the specific person who'd search for this (a job title, a hobby, an industry), not just "everyone"?
- When I search this exact phrase on a marketplace, is the result count reasonable — not empty (no demand), not overwhelming (no differentiation)?
- Can I actually produce a convincing worked example in this niche today, or would I be faking expertise?
- Is there a real recurring task here, or a one-time novelty someone would only buy once ever?

Novelty niches burn out fast

A prompt built around a viral trend or a one-off meme format can spike in sales briefly, then die completely once the trend passes. A niche tied to a genuine, ongoing task (a job function, a recurring content need) keeps selling long after the initial launch.

Coming up next

Chapter 5 covers the actual workflow for writing and testing a prompt properly once you've picked a niche — the step Chapter 3's self-check assumes you've already done.

CHAPTER 5: WRITING AND TESTING PROMPTS

Selling AI Prompts

Course 1 · Chapter 5 · A Systematic Workflow for Writing and Testing Prompts

Chapter 3's self-check assumed a prompt had already been through a real testing process. This is that process — a repeatable sequence that turns a first draft into something you can actually stand behind, applied consistently to every prompt you build rather than done once and forgotten.

The Workflow

1 Start from a specific goal, not a vibe

Write down exactly what output you want, for whom, before typing a single word of the prompt itself. "A good image" isn't a goal; "a consistent, warm-lit product photo style for handmade jewellery listings" is.

2 Draft v1, then run it many times

For image prompts, run it across several different seeds. For text prompts, run it against several different realistic inputs (a different product, a different name, a different scenario each time). One good output proves nothing — consistency across many runs is the actual test.

3 Log where it breaks

Note every run that went off-topic, ignored an instruction, or produced something unusable. These failure modes — not the successes — are what tell you what to fix next.

4 Add one guardrail at a time

For each failure mode found in Step 3, add a specific constraint that closes it — per Chapter 3's own guardrails pillar. Change one thing at a time rather than

rewriting the whole prompt at once, or you won't know which change actually fixed the problem.

5 Re-test after every change

Repeat Step 2 after each guardrail added. A fix for one failure mode can quietly introduce a new one — the only way to catch that is testing again, not assuming the fix worked.

6 Document the final version

Record the exact final wording, the model/tool and version it was tested against, and any known limits (input types it doesn't handle well, settings that matter). This is what Chapter 3's documentation pillar actually produces.

7 File it into a personal prompt library

Keep every finished, tested prompt organised by niche and category in one place (a spreadsheet or a document works fine to start). This is what Chapter 6 draws on for packaging, and what eventually lets you bundle related prompts into themed packs rather than starting from zero each time.

Testing Differs by Category

Image Prompts

- Vary the seed, keep the prompt fixed
- Check consistency of style, composition, and quality
- Model checkpoints are usually stable once released — less ongoing drift risk

Text / Chat Prompts

- Vary the input (different names, topics, scenarios), keep the prompt fixed
- Check consistency of length, tone, and factual accuracy
- Underlying chat models update more frequently — higher ongoing drift risk

PROMPT DRIFT IS REAL

A text prompt that worked perfectly against one model version can behave differently after the provider quietly updates that model — the exact wording never changed, but the model interpreting it did. Re-test listed text prompts periodically rather than assuming a prompt that sold well last month still performs identically today. Image prompts are comparatively more stable once a model checkpoint is released.

Coming up next

Chapter 6 covers turning a tested, documented prompt from your library into an actual listing — sample outputs, before/after comparisons, and bundling related prompts into packs.

CHAPTER 6: PACKAGING YOUR PROMPTS FOR SALE

Selling AI Prompts

Course 1 · Chapter 6 · Packaging Your Prompts for Sale

A tested, documented prompt from your Chapter 5 library still isn't a listing. Packaging is the layer that actually converts a browsing visitor into a buyer — and it matters just as much as the prompt itself, since a buyer has to trust the result before they'll ever type it into an AI tool themselves.

Anatomy of a Strong Listing

Sample outputs

For image prompts: several example images showing the range and consistency, not just one best-case shot. For text prompts: a full worked before/after example, since there's no single picture to sell it on (Chapter 2).

A title that names the outcome, not the mechanism

"Property listing descriptions that highlight neighbourhood features" sells; "a real estate prompt" doesn't. Buyers search for what they want to happen, not for the word "prompt."

What's included, plainly stated

The exact prompt text, any accompanying instructions, and whether it's one prompt or a set — no ambiguity about what the buyer receives.

Requirements and limitations

Which tool and version it was tested against (from Chapter 5's documentation step), and anything it doesn't handle well. Buyers who know the limits upfront leave better reviews than ones who discover them the hard way.

Usage instructions

How to fill in the placeholders from Chapter 3, in plain language — assume the buyer has never written a prompt before.

Single Prompts vs. Themed Packs

Single Prompt

- Lower price point, lower barrier to a first sale
- Good for testing whether a niche has real demand at all
- Each new listing starts its discovery/reviews from zero

Themed Pack

- Higher perceived value — "10 prompts for X" reads as a complete solution, not a single tool
- Higher average order value per sale, same amount of marketing effort
- Draws directly on the personal prompt library built in Chapter 5, so later packs get cheaper to assemble

A common pattern: launch with a few strong single prompts in a chosen niche (Chapter 4) to confirm real demand, then bundle the best-performing ones into a themed pack once you have several that reliably sell.

Don't showcase what isn't reproducible

A cherry-picked, unusually good output that the actual sold prompt can't reliably reproduce is the single fastest way to earn refund requests and bad reviews. Every sample shown in a listing should be a genuinely typical result from the testing done in Chapter 5 — not the best output out of fifty attempts.

Coming up next

Chapter 7 compares where to actually list this — dedicated prompt marketplaces vs. general platforms like Gumroad or Etsy vs. your own storefront.

CHAPTER 7: WHERE TO SELL

Selling AI Prompts

Course 1 · Chapter 7 · Where to Sell: Marketplaces Compared

A packaged listing from Chapter 6 needs somewhere to actually be found. The three broad options trade off discovery against commission in genuinely different ways — and most sellers who stick with this end up using more than one at once.

The Three Options

OPTION	BEST FOR	TYPICAL FEE	DISCOVERY
Dedicated prompt marketplaces	Getting your first sales fast in a niche buyers are already searching for	~20% commission	Built-in, but shared with every other prompt seller
General digital-goods platforms (Gumroad, Etsy)	Themed packs, a nicer storefront presentation, broader (not prompt-specific) audience	Lower/flat fee, varies by platform	Weaker built-in discovery for "prompts" specifically — needs some of your own traffic
Your own storefront/website	Repeat buyers, full pricing/branding control, keeping 100% after payment processing	Payment processor fee only (typically 2–3%)	None built-in — every visitor has to be driven there yourself

How These Actually Combine

Discovery and commission are directly in tension: a dedicated marketplace hands you an audience already searching for exactly what you sell, at the cost of a real cut of every sale, forever. Your own storefront keeps nearly all the money, but has zero audience of its own on day one.

A COMMON MULTI-CHANNEL PATTERN

List on a dedicated marketplace first, specifically for discovery — it puts your work in front of buyers who are already looking. Once a buyer finds you there and likes your work, direct repeat business (a themed-pack bundle, a custom request) to your own storefront or a general platform, where the marketplace isn't taking a cut of every future sale from that same relationship.

Read the actual terms before listing

Fees, approval requirements, and rules around what counts as an acceptable prompt listing all vary by platform and change over time — treat the figures above as a starting comparison, not a locked-in reference, and check the current terms directly on whichever platform you choose before listing anything.

Coming up next

Chapter 8 covers actually pricing what you list — single-prompt vs. bundle vs. subscription pricing, and what genuinely justifies charging a premium.

CHAPTER 8: PRICING YOUR PROMPTS

Selling AI Prompts

Course 1 · Chapter 8 · Pricing Your Prompts

Prompts feel easy to copy, which creates constant downward pressure on price. But a prompt that survived Chapter 5's real testing process — documented, guardrailed, niche-specific — is a genuinely different product from a one-line guess, and deserves pricing that reflects that.

Three Pricing Models

Single Prompt

- Lowest individual price, easiest impulse buy
- Good for testing genuine demand in a new niche (Chapter 4)
- Each sale requires its own discovery — no compounding effect

Themed Bundle

- Priced above any single prompt, but at a discount per item
- **Highest average order value for the same marketing effort**
- Draws directly on the prompt library from Chapter 5

Subscription / Catalog Access

- Recurring revenue instead of one-off sales
- Only sustainable if you can genuinely keep adding new prompts regularly
- Best suited to sellers who've already built a large, active library

What Actually Justifies a Premium

Four things separate a premium-priced listing from a race-to-the-bottom one, and all four trace back to earlier chapters: **documented reliability** from real multi-run testing (Chapter 5), **genuine niche specificity** rather than a generic category (Chapter 4), **workflow complexity** for multi-step prompts that solve a whole process rather than one output (Chapter 2), and **clear usage instructions** that remove friction for a non-technical buyer (Chapter 6). None of these are cosmetic — each one is real work a buyer is paying to skip.

Rough Pricing Reference

CATEGORY	SINGLE PROMPT	THEMED PACK (8–15 PROMPTS)
Image-generation	\$1.99–\$4.99	\$12–\$25
Text / chat	\$3–\$9.99	\$20–\$40
Workflow / chained	\$9–\$40+	\$40–\$100+

- Does this price feel low enough to trust, but high enough to signal real testing went into it?
- Would I rather sell 3 at a premium price than 30 at a discount price, given the review/support burden?
- Am I pricing based on the effort a buyer skips, or just copying a competitor's number?

Undercutting doesn't build a sustainable niche

Pricing at the very bottom of a category to win on price alone attracts the least loyal buyers and trains your own audience to expect discounts. A premium, well-documented listing in a right-sized niche (Chapter 4) consistently outperforms a race-to-the-bottom one in a crowded category — the same lesson applies here as it does to underpricing freelance services generally.

Coming up next

Chapter 9 covers marketing your catalog, building buyer trust, and the realistic limits of protecting a prompt from being copied once it's sold.

CHAPTER 9: MARKETING, TRUST AND PROTECTING YOUR WORK

Selling AI Prompts

Course 1 · Chapter 9 · Marketing, Trust & Protecting Your Work

A well-priced, well-packaged listing (Chapters 6 and 8) still needs to actually be found, and needs a buyer to trust it enough to pay before ever seeing the real output. This chapter covers building that visibility and trust — and gives an honest answer to the question every new seller eventually asks: what actually stops someone from just copying what you sold them?

Building Visibility

A visible output portfolio

Share sample outputs (not necessarily the exact prompt text) on social platforms relevant to your niche — the same output-first proof that sells a listing also markets it.

Community presence

The same forums and communities used for demand research in Chapter 4 double as a marketing channel — genuinely helpful answers build reputation long before any sales pitch does.

Encouraged reviews

Ask buyers directly for a review after a successful use. Reviews are the single strongest trust signal a new visitor sees before ever trying the prompt themselves.

Responsive support

Helping a buyer who's stuck on a placeholder (Chapter 6's usage instructions) turns a one-off purchase into a repeat customer and a positive review.

The Honest Answer on Protecting Your Work

Copyright protection for prompts is genuinely thin. Short text strings and functional instructions often don't meet the originality threshold copyright law expects, and this

varies by jurisdiction. Don't assume a prompt is legally protected the way a full creative work would be.

Platform terms are the main practical recourse. Most marketplaces prohibit reselling a purchased listing as-is and offer a takedown process for exact duplicates — that's a platform-enforced rule, not a copyright claim.

Literal resale of your exact prompt is a smaller risk than it feels like. Most buyers of a \$3–\$10 prompt aren't going to resell it. The more realistic risk is a competing seller independently reaching a similar idea and reworded prompt — which no amount of legal protection stops anyway.

YOUR REAL PROTECTION IS REPUTATION, NOT LAW

A continuously updated library (Chapter 5), a track record of reliable results, and genuine buyer trust are a moving target competitors can't simply copy by lifting your prompt text. That's a stronger long-term defense than any legal protection a short piece of text is likely to have.

This isn't legal advice

Copyright and platform-terms specifics vary by jurisdiction and by platform, and both change over time. Treat this chapter as a realistic starting orientation, not a substitute for checking a platform's actual current terms or consulting a professional if a real dispute ever comes up.

Coming up next

Chapter 10 — the final chapter — turns everything from Chapters 1–9 into a concrete 30-day action plan, from your first tested prompt to your first marketplace sale.

CHAPTER 10: YOUR 30-DAY ACTION PLAN

Selling AI Prompts

Course 1 · Chapter 10 · Your 30-Day Action Plan

Everything from Chapters 1–9 — the categories, testing, packaging, platforms, pricing, and trust — comes together here as one concrete sequence. Thirty days is genuinely enough for a first sale in this model, since the core costs are time and iteration rather than money or a large audience.

1–7**Days 1–7: Pick a niche and build your first prompts**

Run the niche-selection checklist from Chapter 4, then draft and test 3–5 prompts in that niche using Chapter 5's workflow — multiple runs, logged failure modes, guardrails added one at a time — until each one passes Chapter 3's self-check.

8–10**Days 8–10: Package your first listings**

Take your best 3 tested prompts and build real listings around them — sample outputs, a title naming the outcome, requirements, and usage instructions — following Chapter 6's listing anatomy.

11–14**Days 11–14: Choose a platform and go live**

Pick a primary platform from Chapter 7's comparison (a dedicated marketplace is the usual starting choice for discovery), price each listing using Chapter 8's reference table, and publish your first 3 single prompts.

15–21**Days 15–21: Build visibility and launch a themed pack**

Start sharing sample outputs in the communities identified back in Chapter 4, and once your individual prompts show real demand signals (views, favorites, or early sales), bundle your strongest performers into a themed pack per Chapter 6's bundling guidance, priced per Chapter 8.

22–
25**Days 22–25: Collect reviews and support early buyers**

Ask satisfied buyers for a review directly, and respond quickly to anyone stuck using a prompt — per Chapter 9, this is what actually starts compounding trust rather than any single listing on its own.

26–
30**Days 26–30: Review, expand, and re-test**

Look honestly at what sold and what didn't. Add 3–5 more tested prompts to your library in whichever niche performed best, and re-run Chapter 5's testing workflow on your earliest listings to catch any prompt drift before deciding whether to add a second platform.

THE MOST IMPORTANT THING TO DO TODAY

Pick one small, specific niche idea and run it through Chapter 5's testing workflow today — even if it isn't the perfect niche from Chapter 4's checklist. A single tested, documented prompt in hand teaches you more about this business than another hour spent researching which niche might be best.

Course complete

That closes all ten chapters — the opportunity and realistic income expectations (Ch.1), the three prompt categories (Ch.2), what makes a prompt genuinely sellable (Ch.3), niche selection (Ch.4), the write-and-test workflow (Ch.5), packaging and bundling (Ch.6), where to sell (Ch.7), pricing (Ch.8), marketing and the honest limits of protecting your work (Ch.9), and this 30-day plan tying it all together. A combined, portable PDF of the full course is available for offline reference.

This course is intended for educational and exploration purposes. Fees, price points, and platform terms represent typical ranges at the time of writing and vary by platform, niche, and individual circumstances. Always verify current terms directly on whichever platform you actually use before pricing or listing anything.