The AI-Powered Amazon Seller

How to Use Artificial Intelligence to Find Products, Write Listings, Run Ads, and Scale Your Brand — From First Sale to Eight Figures

by Kevin James Gill

Copyright and Disclaimer

This book is educational material. It is not legal, tax, accounting, or financial advice. Consult a qualified professional before making decisions about your business entity, taxes, contracts, product liability, or intellectual property.

The author is not affiliated with, endorsed by, or sponsored by Amazon.com, Inc. or any of its subsidiaries. "Amazon", "Amazon FBA", "Seller Central", "A+ Content", "Rufus", "Alexa", "Alexa for Shopping", and related names and logos are trademarks of Amazon.com, Inc. Every software product, service, and platform named in this book is the property of its respective owner, and nothing here implies endorsement in either direction.

Pricing, product features, and Amazon policy change constantly. Anything described in this book may be outdated by the time you read it. Verify every policy-sensitive action against current Seller Central policy and Amazon's published terms before you act on it. You are responsible for your own account.

Some links in this book may be affiliate links, meaning the author may earn a commission at no additional cost to you.

Results are not guaranteed. No income claim is made or implied. Selling on Amazon involves financial risk, including the risk of losing your capital.

How to Use This Book

This book has five parts, thirteen chapters, and four appendices. It moves in the order the work actually happens: understanding what AI is genuinely good at, choosing your tooling, finding and validating products, building listings and creative, running ads and pricing, then operations, brand, and scale.

Read it in order the first time. After that, treat it as a reference.

The two boxes. Every chapter carries two recurring boxes.

STARTING OUT is for a seller working toward or just past their first sales. Small budget, no team, doing everything personally. The advice in these boxes assumes your constraint is money and time, not headcount.

SCALING ($500K+) is for an operator with a real catalogue, staff or contractors, and enough data volume that decisions become statistical rather than intuitive. The constraint here is coordination and consistency, not effort.

Read your own box carefully. Skim the other one anyway. Sellers starting out should know what the machinery looks like on the other side, because you are building toward it and the decisions you make now — how you name SKUs, how you store your data, whether you document anything — determine how painful that transition is. Operators at scale should read the STARTING OUT boxes because you will launch new brands, hire people who are at that stage, and forget how much of your process is muscle memory rather than written procedure.

Do This Today. Every chapter ends with a single action that takes about twenty minutes. Not a summary. Not a reflection question. One concrete thing you do in your account, your spreadsheet, or your chat model window before you move to the next chapter.

The book is worthless if you skip them. Reading about AI-assisted keyword work produces nothing. Doing twenty minutes of it produces a file you will use next week. Thirteen chapters, thirteen actions, roughly four hours of total work — and at the end you have a functioning system rather than a set of opinions.

Why there are almost no tool names. Specific products appear in exactly two places: Chapter 2, where we walk the tooling landscape, and Appendix B, where the current stack is listed. Everywhere else, this book talks about categories — a repricing tool, a bid-management platform, a general-purpose chat model, a review-analysis tool.

That is deliberate. Named software goes stale in months. Products get acquired, change pricing, break their APIs, or quietly get worse. Mechanisms do not. The logic of how to structure a keyword harvest, how to decide when a bid is genuinely underperforming, how to read a review corpus for product defects — that survives every tool cycle.

So you learn the job first and pick the software second. When a tool you rely on disappears, you will already understand what it was doing for you, and replacing it becomes a shopping trip instead of a crisis. It also means this book gets revised in two places instead of eighteen.

Two appendices to keep open while you work. Appendix A is fifty copy-paste prompts, organised by task. Appendix C is 30/60/90-day checklists. Those two are working documents. The rest is reading.

About the Author

Kevin James Gill is an operator. He works with sellers on the unglamorous parts of the business — catalogue hygiene, advertising structure, forecasting, and the process work that decides whether a brand survives its own growth — and builds automation for the tasks that should never have been done by hand.

He writes about applied AI: what it actually does well, what it fails at in ways that cost money, and how to tell the difference before you have bet a quarter on it.

He has no interest in guru theatre. No rented cars, no screenshots of dashboards, no promise that this is easy. Selling on Amazon is a real business with real capital risk and a platform that changes the rules without asking. AI is a genuine advantage inside that business, and it is not a shortcut around it.

If this book saves you a month of avoidable mistakes, it did its job.

Introduction: The Year the Work Changed

For most of Amazon's history, selling on it was a labor problem disguised as a business problem.

You wanted to know what to sell, so you spent nights scrolling categories, exporting keyword lists, and squinting at sales-rank estimates. You wanted a listing that converted, so you read a hundred reviews with a notepad open, then wrote five bullets and rewrote them twelve times. You wanted profitable ads, so you pulled a search-term report on Sunday, sorted by ACOS, and hand-adjusted bids until your eyes hurt. You wanted to not run out of stock, so you built a spreadsheet with a lead-time column and a fudge factor and hoped.

That was the job. Not the strategy — the strategy was a few decisions a year. The job was the grinding, mechanical middle: gathering, reading, sorting, drafting, adjusting.

That middle is now the part a model does in minutes.

Hand a general-purpose chat model four hundred reviews and ask it what buyers complain about, and you get a ranked list of failure modes in the time it takes to refill your coffee. Give it your category, your constraints, and your competitors' bullets, and it drafts twenty listing variations before you've finished reading the first one. Point it at a search-term export and it will tell you which terms are bleeding you, which are starving, and which look like a different product entirely.

None of that used to be fast. Now it is. Which means the thing that separates sellers is no longer who is willing to grind — everybody's grind is now cheap — but who has the judgment to point the grinding at the right thing, and the discipline to check the output before it goes live.

That is what this book teaches.

What The Machine Actually Took Off Your Plate

Be specific about this, because vague enthusiasm about AI is how people waste six months.

What models are genuinely good at, in this business, is a short list with a common shape: tasks where the input is a large pile of language or numbers, the output is a summary, draft, or classification, and a human can verify the result quickly.

Reading reviews and pulling out the recurring complaint. Turning a messy keyword export into themed clusters. Drafting bullets in five different angles so you can test them instead of agonizing over one. Writing the first version of a negative-feedback response. Translating your listing into a market you don't speak the language of. Explaining what a Business Report column actually means. Turning a page of raw ad data into three sentences about where money is going.

The common thread is that these are all *first drafts and first passes*. The model gets you to eighty percent in two minutes on work that used to take two hours. You supply the last twenty percent, which is where all the value was hiding anyway.

That is the whole trade. And it is a very good trade — but only if you actually do the last twenty percent.

What It Does Not Do, And Will Not Do For You

Now the honest part, because this is where sellers get hurt.

It does not pick a good product for you. A model can summarize demand signals, cluster niches, and argue both sides of a category. It cannot tell you whether you can source that item at a landed cost that leaves you a margin, whether your supplier will still be answering emails in March, whether the category is about to be flooded, or whether you have any real advantage in it. Product selection is a capital-allocation decision made with incomplete information. That is judgment. Nobody has automated it, and the people telling you otherwise are selling something.

It does not know your margins. It does not know your unit cost, your freight, your storage bill, your return rate, your referral fee, your PPC spend per unit, or your cash conversion cycle. Unless you feed it those numbers, anything it says about profitability is a guess dressed as an answer. Feed it your real numbers and it becomes a useful calculator. Don't, and it becomes a confident liar.

It hallucinates numbers. This one deserves blunt treatment. Ask a chat model for a category's average conversion rate, or a competitor's monthly units, or "the typical ACOS in home goods," and it will very often produce a clean, specific, plausible figure with no source behind it. It is not lying on purpose — it is predicting the shape of an answer. Treat every number a model produces as a claim requiring a source. If the number came from data you supplied, verify the arithmetic. If it came from nowhere, it *is* from nowhere. Throughout this book, when a number matters, I will point you at where to pull it in your own account rather than quoting a statistic at you.

It will happily write you a listing that gets your account suspended. This is the big one. A model trained on the whole internet has absorbed a lot of marketing language, and marketing language is full of claims that are fine on a Shopify store and catastrophic on Amazon. Medical and health claims. "FDA approved." "Best seller." "Number one rated." Competitor callouts. Guarantee language that conflicts with policy. Safety claims about materials you have not tested. Keyword stuffing that reads fine to a model and looks like manipulation to an enforcement system. The model has no idea which of these will trigger a listing takedown, a suppression, or a policy strike, because it has no idea what your account looks like or what got enforced last quarter.

Say it plainly: **every claim in your listing is your legal responsibility, not the model's.** If a generated bullet says your supplement supports immune function, you made that claim. If a generated A+ module says your product is certified, you asserted a certification. Regulators and Amazon's enforcement systems do not care that a language model drafted it. There is no "the AI wrote it" defense — not with Amazon, not with the FTC, not with a plaintiff's attorney.

The rule for the rest of this book: **AI drafts, you approve.** Never the reverse.

Who This Is For

This book assumes you are already in the arena, or about to walk into it.

You have at least one live listing. You have seen a real order come in, paid a real referral fee, and had a real return. Or you are days away from launch, with a product decided and inventory moving. You know what a search-term report is even if you don't enjoy reading one. You have some money at risk.

That matters because almost everything in here is about applying leverage to an existing operation. You cannot automate a business you do not yet understand. If you have never launched anything, the chapters on research and validation will still help you — but the ads, inventory, and analytics chapters will read like a foreign language until you have data of your own to point them at.

Who this is not for: anyone looking for a passive income button.

There isn't one. There has never been one, and AI did not create one. What AI created is a much higher ceiling for people who are willing to run the operation. If your plan is to have a model pick a product, write a listing, launch some ads, and wake up to deposits, you will lose money faster than sellers did before — because the tools now let you execute a bad decision at speed and scale. That is the actual risk of this era. Not that AI does the work badly, but that it does the wrong work extremely well.

There is no shortcut in this book. There is a great deal of leverage.

How The Chapters Are Built

Every chapter runs the same way, so you can move fast once you know the shape.

The teaching is in the body text. Alongside it you will find two recurring boxes. The first is a **prompt box** — the actual working language to give a model, written so you can adapt it rather than copy it blindly, along with what to feed it and what a good output looks like. The second is a **guardrail box** — the specific way that chapter's technique goes wrong: the policy tripwire, the hallucination pattern, the automation that quietly burns budget while you sleep.

Each chapter closes with a single **Do This Today** action. One task, sized to be finished in one sitting, using your own account and your own data. It is there because reading about leverage produces nothing. Doing one small thing produces a result you can look at, which produces the next thing. Work the actions and the book compounds. Skip them and it is entertainment.

That is the whole structure. The How to Use This Book section covers the rest.

The Expectation To Set Before You Start

Here is the honest framing, and I would rather you hear it now than feel misled in Chapter 8.

This is a leverage book, not a shortcut book. Nothing in here reduces the amount of thinking you have to do. It reduces the amount of typing, reading, sorting, and drafting — which frees up your thinking for the decisions that actually move money.

The practical shape of that change is throughput. Before, testing a new set of bullets was a half-day project, so you did it when something was clearly broken. Testing five image concepts meant a photographer and two weeks. Rebuilding a campaign structure meant a weekend. So you ran one experiment a month, and each one carried a lot of emotional weight because it cost so much to run.

Now the drafting cost is close to zero. The seller who wins in this environment is not the one with the cleverest prompt. It is the one who runs ten experiments a week instead of one, reads the results honestly, kills what fails without ego, and keeps the two things that worked.

That is the entire game. Cheap experiments, ruthless reading of results, fast iteration. AI is the thing that made the experiments cheap. You still have to run them, and you still have to be honest about what the data says — which, in my experience, is the harder of the two skills.

One more expectation: this is a moving target. Amazon's policies on review solicitation, automated pricing, generated imagery, and account automation change, and they change without a press release. Every time this book touches one of those areas, it tells you what the risk is and why — and it tells you to go read the current policy in Seller Central before you act. Do that. A book cannot be your compliance department.

Where We Start

Chapter 1 is about the ground shifting under everyone at once. Buyers are searching differently now that they can ask conversational questions instead of typing keyword fragments. Competitors have the same drafting speed you do, which means listing quality is inflating across every category. And the platform itself is running more of its surface — search, recommendations, ad placement — through models rather than rules.

You need to understand that landscape before you start applying tools to it, because half the tactics that worked in 2019 now work against you.

Turn the page.

Chapter 1: Why Amazon Changed Under Your Feet

There's a specific kind of seller I keep meeting. They built a real business between 2016 and 2022. They learned the game properly — reverse-engineer the competition, harvest every keyword variant, cram them into the title and the backend fields, drive external traffic, win the badge, defend the rank. It worked. It worked so well that they stopped questioning whether it was still the game.

Then somewhere in the last couple of years, the ground moved. Their listings didn't get worse. Their conversion rate didn't crater overnight. But new products stopped taking off the way they used to. Ad spend bought less. Competitors they'd never heard of started showing up in placements they thought they owned. And when they asked why, everybody gave them the same non-answer: "the algorithm changed."

The algorithm did not change. The *question the algorithm is answering* changed. That's a much bigger deal, and it's what this chapter is about.

From Matching Words to Matching Meaning

For most of Amazon's history, search was a matching problem. A shopper typed a string. Amazon looked for products whose indexed text contained that string or something close to it, ranked the matches by a blend of relevance and sales performance, and served a page. Your job as a seller was to make sure that whatever string the shopper typed, your product's text contained it.

That's why the entire optimization industry grew up around keyword coverage. Get indexed for more terms, get shown for more searches, get more sales. The craft was in the harvesting and the placement — which terms go in the title, which go in bullets, which go in the backend search fields where the shopper can't see them. Nobody was writing for humans. We were writing for a text-matching index and hoping humans could tolerate the result.

Amazon has since layered semantic understanding on top of that index. Their internal work on this is publicly discussed under the name COSMO — a model built to understand *why* someone is shopping, not just what letters they typed. And the customer-facing half is the shopping assistant embedded in the Amazon app and site, which lets a shopper ask a question in plain language and get an answer with products attached. It launched as Rufus, and on 13 May 2026 Amazon merged it with Alexa+ and renamed the combined product **Alexa for Shopping**. You will still see the old name everywhere in seller forums and in course material written before that date. The name is the only thing that changed in a way that matters to you — it is the same shift from matching text to answering questions, and the plumbing underneath is the same. Expect the branding to keep moving; the behaviour is what you optimise for.

The mechanism underneath both is worth understanding, because it explains nearly every tactical change that follows. Modern search doesn't compare your text to the query as strings. It converts both into mathematical representations of meaning — vectors in a high-dimensional space — and measures how close they sit. Two phrases that share no words at all can land right next to each other if they mean the same thing. Two phrases that share every word can land far apart if the intent behind them differs.

That's a genuine break with how the old index worked, and it's the load-bearing idea for the rest of this book. If the underlying math is unfamiliar and you want it properly explained rather than gestured at, the concepts are their own subject and we cover them plainly in the Semantic Search and Embedding Space entries in the glossary.

Here's what it means in practice. A shopper who searches "gift for my dad who has every tool already" is not typing a keyword. There is no keyword. There's a situation, a relationship, a constraint, and an implied budget. Old search would have choked on that string or matched the word "tool" and served you a hammer. New search understands that this is a gifting query, that the recipient is experienced, that novelty and quality matter more than price, and that the answer is probably not a hammer.

Your listing does not need to contain the phrase "gift for my dad who has every tool already." It needs to *be about* something that satisfies that intent, in language clear enough that a machine reading it can tell.

Why Buyers Started Asking Instead of Typing

Search behaviour follows the interface. When the interface was a text box that punished long queries with bad results, shoppers learned to type short, blunt terms — "yoga mat," "yoga mat thick," "yoga mat thick non slip." That wasn't natural language. That

was people learning to speak machine.

Give those same people an assistant that answers questions, and within a few sessions they stop compressing. They start asking things like "which yoga mat won't slip on carpet" and "is this thick enough for someone with bad knees." Not because they got smarter, but because the tool stopped punishing them for being specific.

That behavioural shift is the whole story, and it has three consequences you'll feel in your numbers before you can name them.

The search tail got longer and stranger. More queries, each with fewer impressions, phrased in ways your keyword tool never suggested because no tool suggests a sentence. Your search-term report is drifting toward a shape that looks like noise if you're still reading it as a keyword list.

The question layer sits above the product layer. A shopper who asks the assistant a question may get a synthesized answer that mentions considerations, trade-offs, and product characteristics before they ever see a grid of listings. You are now competing to be *cited in the answer*, not just ranked in the grid. Nobody has a reliable playbook for that yet, including me. What we do know is that clear, factual, specific product information is what gets pulled into answers, and vague marketing language isn't.

Objection-handling moved earlier. Buyers used to hit objections on your detail page, in the reviews, in the Q&A. Now they raise those objections to an assistant *before* they reach you. If your listing doesn't contain the answer, something else does, and that something else is a competitor.

Tip: Go read the "Customer Questions" and the one-, two-, and three-star reviews on your top three competitors. Not for schadenfreude — for the exact phrasings of what worries buyers in your category. Those sentences are, almost verbatim, the questions shoppers are now asking an assistant. Every one of them that your listing answers explicitly is a place you can be surfaced and they can't. This is the single highest-return hour of unglamorous work in this entire book, and almost nobody does it because it isn't automatable and doesn't feel like strategy.

Why Keyword Stuffing Now Costs You Money

This is the part that makes long-time sellers uncomfortable, so let me be direct: the thing that used to be your edge is now a liability. Not neutral. Negative.

Three mechanisms, and they compound.

One: dilution. When meaning is computed from your whole text rather than matched term by term, every irrelevant word you add drags your product's representation away from what it actually is. A title stuffed with fourteen loosely-related terms produces a blurry semantic signal — a product that's vaguely near a lot of things and precisely near nothing. A clean, specific title lands your product firmly in the neighbourhood where its real buyers are looking. Under the old index, extra terms were free lottery tickets. Under this one, they're tax.

Two: intent mismatch punishes you twice. If stuffing does get you shown for a query you don't genuinely serve, you get an impression from someone who won't buy. Low click-through, or a click and a bounce. Amazon reads that as evidence you're not a good answer for that query and shows you less — for that query and, over time, for its neighbours. You paid for the impression with your own relevance score.

Three: humans. Stuffed copy reads like a machine wrote it for another machine, and shoppers can tell. A shopper who reaches your page after asking a careful question and finds a wall of comma-separated noun phrases has just learned something about how much care went into the product.

The failure mode I see most often isn't crude stuffing anyway — most sellers cleaned that up years ago. It's the sophisticated version. A well-formatted listing with five tidy bullets, each front-loaded with a keyword phrase in capitals, each saying essentially nothing. "PREMIUM QUALITY MATERIALS — Made with premium quality materials for lasting durability." That bullet is a keyword with a period after it. It answers no question, resolves no objection, and gives a semantic model nothing to work with. It looks professional and it performs like filler.

What replaces it is specificity. Not "premium materials" but what the material is, what it does, and what it means for the buyer. Not "perfect for any occasion" but the two or three actual occasions. Specific text is what a meaning-based system can actually

place, and it's what a human actually wants to read. That's the rare case where the machine's interest and the customer's interest point the same direction — take the win.

> **STARTING OUT** — Audit your category the way a shopper would ask, not the way a seller would search.
>
> This costs nothing and takes an afternoon. Open the Amazon app on your phone as a normal shopper — not Seller Central, not a keyword tool. Find the shopping assistant and ask it the ten questions a real buyer in your category would ask. Not "yoga mat." Ask "what thickness of yoga mat do I need if I have knee problems," "is a cork mat better than rubber," "which mats don't smell like chemicals when new."
>
> Write down, in a plain document, three things for each question: the answer you got, the products it surfaced, and — this is the important one — whether *your* product could have been the answer and just wasn't.
>
> Then take the same ten questions to any free general-purpose chat model and ask it to answer them as if it were a knowledgeable shopping advisor. You're not looking for accuracy. You're looking at what *attributes* it reaches for when reasoning about your category. Those attributes — thickness, material, odour, grip on carpet — are the dimensions along which your product is now being understood. If your listing doesn't state where you sit on each one, you're invisible on that dimension.
>
> Total cost: nothing. Total time: an afternoon. You will end it knowing more about your category's real demand than a paid keyword tool would have told you.

The Three Places AI Touches Your Business

Everything in the rest of this book fits into three buckets. I want them named early, because sellers routinely over-invest in one and ignore the other two, and the one they over-invest in is almost always the wrong one.

What you sell. Product selection, category assessment, differentiation, and the read on whether demand is real. This is where AI is most useful and most under-used. The work is synthesis: pulling apart hundreds of competitor reviews to find the complaint that shows up in every single one, mapping which sub-segments of a category are genuinely underserved versus merely crowded, pressure-testing a product idea against the objections a buyer would actually raise. A language model is very good at reading a thousand reviews and telling you what they have in common. A human is very good at deciding what to do about it. The mistake is skipping the human half.

What AI cannot do here is tell you whether a product will sell. It has no access to demand it hasn't seen, it will happily invent a market that doesn't exist, and it has a strong bias toward telling you your idea is interesting. Treat it as a research analyst with an unlimited reading budget and no judgment whatsoever.

How you list it. Titles, bullets, descriptions, A+ content, backend attributes, images, video. This is where most sellers start, and it's the bucket where AI produces the most output and the least differentiation. Everyone has access to the same models. If your process is "paste competitor bullets, ask for better bullets," you will produce something indistinguishable from what every other seller in your category produced with the same prompt. Generated copy converges. That's a property of how the models work, not a prompting failure you can fix with a cleverer prompt.

The leverage isn't generation, it's *grounding* — feeding the model things it couldn't have guessed. Your actual specifications. Your actual review complaints. The three questions your customer service inbox answers every week. The model turns your proprietary information into good copy. Without that input, it turns thin air into average copy.

One hard line while we're here: generated product imagery. Amazon's policies on image accuracy are strict and enforcement is real. An image that shows the product in a way that doesn't match what ships — invented textures, a different colourway, accessories not included, an AI-rendered scale that's flattering rather than accurate — is a misrepresentation, and misrepresentation is a listing-suppression and account-health issue, not a growth hack. The rules on generated imagery specifically have shifted more than once and will shift again. Read the current image and product-detail-page policy in Seller Central before you publish anything generated, and re-read it before your next launch. Do not take my word or a forum post's word for what's allowed this quarter.