



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

AI & SEARCH

How to Get Your Products Recommended in ChatGPT Shopping and AI Assistants

By Scott Tischler · July 28, 2026 · 8 min read

Ask ChatGPT for "the best carry-on under \$200" and it will not hand you ten blue links. It hands you an answer: a short list of specific products, often with prices, pros, and a reason each one made the cut. That answer is the new storefront. And for a growing share of shoppers, it is the only storefront they will see before they buy.

This is **agentic commerce** — people asking an AI assistant what to buy, and increasingly letting it compare, shortlist, and even check out on their behalf. The shift is not coming. It is already reshaping how demand gets discovered. If your brand is not one of the names the assistant surfaces, you are not losing a ranking. You are losing the sale before the shopper ever knew you existed.

I run Alrecommend.ai, where we help brands become the ones AI systems recommend. Here is how the machine actually decides — and the concrete playbook to earn the pick.

What is AI shopping and how is it different from search?

AI shopping is the process of a person using an assistant like ChatGPT, Gemini, Perplexity, or Copilot to research and select products through conversation instead of browsing. Traditional search gives you a page of options and makes you do the work. An AI assistant does the work and gives you a decision.

That difference is everything. In classic e-commerce SEO, being on page one meant being one of ten choices. In AI shopping, the assistant often names three to five products total — and frequently leads with a single "best overall" pick. The funnel didn't get narrower. It collapsed.

There is no scrolling to position seven. There is no "browse more results." The shopper asked a question and got a shortlist, and the psychology of a shortlist is brutal: the named brands feel vetted, and everything unnamed feels like it doesn't exist.

How do AI assistants decide which products to recommend?

Assistants pull product recommendations from a blend of signals. No single one guarantees inclusion, but the brands that win tend to send strong signals across all of them.

Signal	What it is	Why it matters
Product data feeds	Structured feeds pushed to merchant programs (Google, Bing, retailer APIs)	The assistant's fastest path to accurate price, availability, and specs
Structured product schema	Product, Offer, Review, and AggregateRating markup on your pages	Machine-readable proof of what you sell and how it's rated
Reviews and ratings	Volume, recency, and average score across platforms	Assistants weight social proof heavily when ranking "best"
Retailer and marketplace presence	Being sold and reviewed on Amazon, Walmart, Target, category retailers	Third-party listings are trusted corroboration
Editorial and "best of" coverage	Inclusion in reviews, roundups, and buying guides	Assistants lean on the same sources humans cite
Brand authority	Consistent, corroborated information about your brand across the web	Reduces the model's uncertainty about recommending you

The pattern underneath all of this is simple: **AI assistants recommend what they can verify**. A model is managing risk. It will not confidently name a product it cannot corroborate from multiple independent sources. Your job is to remove every reason for the model to hesitate.

What actually gets a product surfaced?

Three things do most of the heavy lifting.

Corroboration. The assistant wants to see the same claim in more than one place. If your product page says you're the best budget option, that's marketing. If three review sites, a retailer listing, and 4,000 customer ratings say the same thing, that's a fact the model can repeat. Consistency across sources is not a nice-to-have; it is the mechanism.

Structured clarity. Models parse structure faster and more reliably than prose. A page with clean Product schema — name, price, availability, GTIN, review count, average rating — is far easier for an assistant to lift into an answer than a beautiful page that hides its facts in images and marketing copy.

Specificity. "Award-winning quality" is invisible. "4.7 stars across 12,000 reviews, machine-washable, ships in 2 days, \$149" is quotable. Assistants recommend the products they can describe precisely, because precision is what makes an answer feel trustworthy to the human reading it.

The playbook: how to become the recommended pick

Here is what I put in front of brands and retailers who want to win in AI shopping. Work it in order.

1. Fix your product feed and schema first

This is the foundation, and most brands get it wrong. Every product page should carry complete, accurate **Product schema** with price, availability, unique identifiers (GTIN/MPN/SKU), and aggregate rating. Your merchant feeds should be complete and current — no missing prices, no stale availability, no truncated titles. If the assistant can't get clean data, it moves on to a competitor who provides it.

2. Build review depth, not just review scores

Volume and recency matter as much as the average. A product with 8,000 recent reviews at 4.5 stars will usually beat one with 200 reviews at 4.9. Make leaving a review effortless, respond to negative ones publicly, and keep the flow steady so your ratings always look alive. Spread proof across your site, your retailer listings, and independent platforms.

3. Win the third-party corroboration

You cannot be the only source that says you're good. Get listed and reviewed on the major marketplaces in your category. Pursue inclusion in the "best of" roundups and buying guides that assistants draw from — these are the AI era's most valuable placements. Earn genuine editorial coverage. Every independent source that names your product lowers the model's risk of recommending it.

4. Answer the buying question on your own pages

Shoppers ask assistants comparative, use-case questions: "best for small kitchens," "most durable for travel," "good for sensitive skin." Publish content that answers those exact questions clearly and honestly, including where your product fits and where it doesn't. Candor is a feature — assistants and shoppers both reward it, and honest comparison content gets quoted.

5. Make your brand facts consistent everywhere

Same product names, same specs, same claims across your site, feeds, marketplaces, and press. Contradictions create uncertainty, and uncertainty is what stops a model from naming you. Consistency is a ranking signal in disguise.

6. Test it and monitor it like a channel

Ask the assistants directly: "What's the best [your category] under [price]?" Do it across ChatGPT, Gemini, Perplexity, and Copilot. See who gets named and why. This is your new rank tracking. If a

competitor keeps getting picked, reverse-engineer their corroboration and close the gap. **AI recommendation is a channel now — measure it, don't guess at it.**

Why this is urgent, not eventual

The brands establishing strong AI-recommendation signals today are compounding an advantage. Every review, every corroborating source, every clean feed makes the model more confident naming them — and that confidence is sticky. When an assistant learns to trust a product as "the answer" to a buying question, dislodging it is expensive.

The window where this is a competitive edge rather than table stakes is open right now. In agentic commerce, the shelf has room for a handful of names. Decide to be one of them before your category fills up.

The old question was "how do I rank?" The new one is sharper: **when a shopper asks an AI what to buy, does it say your name?** Everything above is how you make sure the answer is yes.

Key takeaways

- ② AI shopping collapses the funnel: assistants name three to five products, so being one of ten choices is no longer good enough.
- ② AI assistants recommend what they can verify — corroboration across independent sources is the core mechanism, not marketing copy.
- ② Clean product feeds and complete Product schema are the foundation; if the assistant can't read your data, it recommends a competitor.
- ② Review depth and recency often matter more than a perfect average score — steady volume signals a product is alive and trusted.
- ② Third-party proof (marketplaces, buying guides, editorial roundups) is the highest-value placement in the AI era because assistants cite the same sources humans do.
- ② Treat AI recommendation as a measurable channel: query the assistants directly, see who gets named, and close the corroboration gap.

Frequently asked questions

How do I get my product recommended by ChatGPT?

Give ChatGPT clean, verifiable signals it can trust: complete Product schema and merchant feeds, deep and recent reviews, presence on major marketplaces, and inclusion in independent buying guides. ChatGPT recommends products it can corroborate from multiple independent sources, so the goal is to make the same accurate facts appear consistently across the web.

What is agentic commerce?

Agentic commerce is when people use an AI assistant to research, compare, shortlist, and sometimes purchase products through conversation rather than manual browsing. The assistant does the shopping work and returns a decision, which means the brands it names capture demand before the shopper ever visits a website.

Does product schema markup actually affect AI recommendations?

Yes. Structured Product, Offer, and Review schema make your price, availability, identifiers, and ratings machine-readable, which lets an assistant lift them directly into an answer. Pages that bury facts in images or marketing prose are much harder to recommend than pages with clean, complete schema.

Are reviews more important than my website for AI shopping?

Reviews are one of the strongest signals because they provide independent corroboration and social proof that assistants weight heavily when ranking "best." Your website still matters for schema and honest comparison content, but review depth, recency, and volume across multiple platforms often decide whether you get named.

How is AI shopping different from traditional e-commerce SEO?

Traditional SEO fights for a spot among ten results on page one; AI shopping fights to be one of only three to five products the assistant names, often with a single "best overall" pick. The funnel is dramatically narrower, so the emphasis shifts from keyword ranking to being the most verifiable, corroborated choice in your category.

Can I pay to be recommended by an AI assistant?

Some assistants are introducing sponsored or shopping placements, but organic recommendations are still driven by verifiable signals — feeds, schema, reviews, and third-party corroboration. Relying on paid placement alone is fragile; the durable advantage comes from being the product the model independently trusts as the best answer.

How do I know if AI assistants are already recommending my products?

Ask them directly. Query ChatGPT, Gemini, Perplexity, and Copilot with the real buying questions shoppers use — "best [category] under [price]" — and note who gets named and why. This is the AI-era equivalent of rank tracking and should be monitored regularly as its own channel.

How long does it take to start getting recommended?

It depends on your starting point, but fixing feeds and schema can improve verifiability within weeks, while building review depth and third-party corroboration is a multi-month effort. The brands that start now compound their advantage, because once an assistant trusts a product as the answer, that position is hard for competitors to dislodge.

About the author. Scott Tischler is the Founder & Chairman of Alrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express, MetLife, and UBS — and executive study at Wharton, Harvard, Yale, and Oxford, he helps businesses become the ones AI recommends.