

Trust Signals Are the New AI Currency

How verified claims, labels, and endorsements drive visibility and ROI in the age of AI shopping

Executive Summary

The future of commerce is here. You're no longer only selling to humans, you're selling to models. AI agents like Google Gemini, ChatGPT, Perplexity, and Claude are already shaping what products people see and consider, and soon they'll be making purchase decisions on behalf of millions of consumers. Humans will consume the outputs, but the critical leverage point is the model. Winning in this new landscape means learning how to convince AI to choose your products.

At Novi, we believe the foundation for better choices is better data. So we conducted robust simulations using real product data across all of today's most popular AI models to understand which data is most critical to making sure AI shoppers can confidently recommend your products.

Our research shows that **trust signals are the strongest determinant of visibility and AI selection in this new world**. Trust signals, which include badges, labels, certifications and endorsements, are the marks of credibility that AI agents look to when deciding which products to recommend.

Our Testing

- Across all major AI shopping models, products with verified trust signals were selected 230-259% more often than random chance, while products without trust signals were selected 86-96% less often than random chance.
- Promotional labels underperformed verified trust signals by 10-58x, and "sponsored" tags actually hurt product selection rate vs. no labels.
- Named authorities (e.g., Ulta Conscious Beauty, MadeSafe, EWG, or Target Clean) beat generic verifiers by 2.0-5.6x.
- Products presenting signals as a badge + optimized text were selected 7x+ more often than those using text descriptors alone.
- Product selection increased +257% from hyper-personalized queries when relevant verification was present and optimized in the product description content.
- When we added verified data to Neutrogena product listings, their products grew selection share from 3% to 9% (+200%), causing CeraVe to lose market share by -5.5 points.

The New AI Shopping Landscape

The Long Tail Opportunity

In Google search, queries average 4-5 words¹. In ChatGPT, queries average 23 words, with some reaching up to 2,717 words². That makes AEO (Answer Engine Optimization) 5+ times larger than SEO. Consumers are asking highly specific questions that never existed in search:

"Help me find fabric softener for my child with eczema in plastic-free packaging."

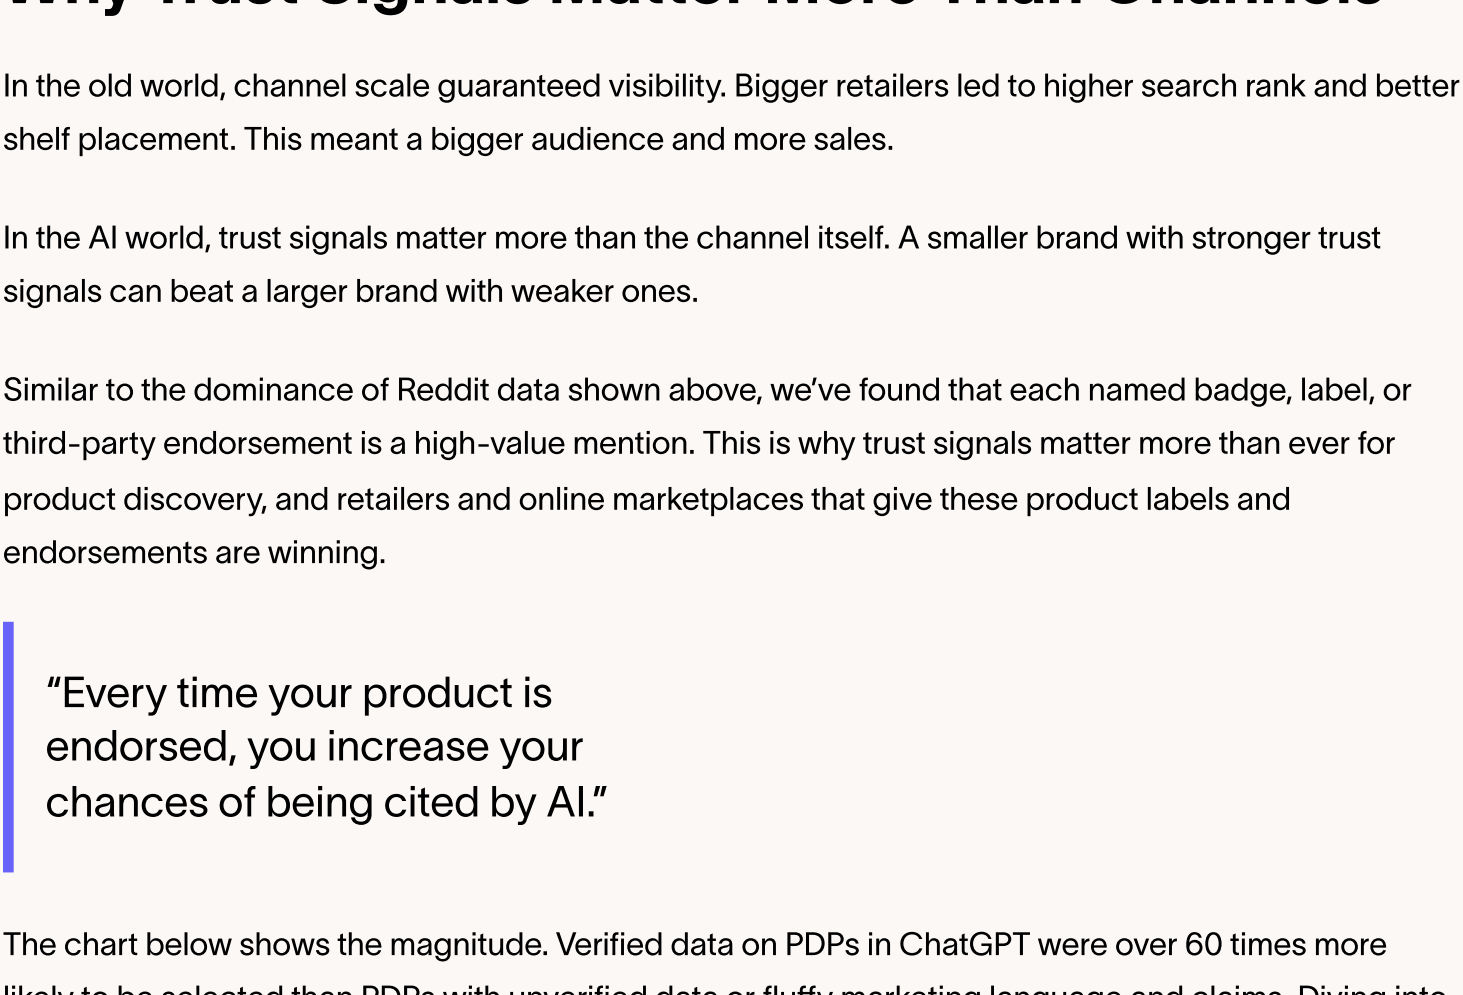
These micro-niche, hyper-personalized queries require structured, credible data. In this world, the question is: Do AI shopping platforms have the information they need to surface your products?

Mentions Beat Rank

In Google and traditional SEO, ranking #1 was everything. In AI, models summarize across multiple sources, and the volume of credible mentions matters more than rank or any single placement.

This marks a profound shift. As one AEO researcher put it: "In AI, being mentioned often beats ranking first."³ But, these mentions must be credible. For example, ChatGPT leans heavily on Reddit as a source because the community polices spam better than any algorithm. A few authentic endorsements across Reddit threads can outweigh a single placement elsewhere.

Figure 1: Top cited sources in ChatGPT Responses



AI Shopping follows the same principle. Each trust signal on a product description page (PDP) plays a similar role to an authentic endorsement in a reddit thread. And our studies show that this is what AI will pick up and trust. The more credible endorsements your product has across the internet, the more likely it is to surface in AI answers because multiple credible mentions carry more weight than a single top-ranked placement.

"Marks of credibility are now the signals AIs listen to. Without them, your products are invisible."

Position Effects in AI Shopping

Unlike human shoppers, who can scroll and compare, AI agents typically choose one product from a fixed set. That means position on the list (the order in which products are presented) can be a unique and powerful factor in AI commerce.

In our ChatGPT (GPT-4o) and Gemini (1.5-pro) experiments, selection was more evenly distributed across the middle positions, with positions 4 and 5 peaking in likelihood for ChatGPT and positions 2 and 6 for Gemini. This suggests that both ChatGPT and Gemini weigh product attributes and trust signals heavily, not just where the product appears.

By contrast, in our Perplexity (Sonar) trials, position effects were much stronger than in GPT-4o, Gemini (1.5-pro), or Claude (sonnet-4). Products in the first position captured 23.5% of selections, more than double the other slots.

This suggests that for ChatGPT and Gemini, verified trust signals and product content are the dominant levers and the model rewards credibility more than placement. For Perplexity and other engines with stronger position bias, brands must optimize both for trust signals and retailer merchandising to maximize visibility.

Position Effects in AI Shopping

For brands: Every incremental trust signal expands your odds of being recommended across AI engines.

For retailers: Trust signal programs and endorsements are no longer optional merchandising levers. They are the price of entry to AI visibility.

Why Trust Signals Matter More Than Channels

In the old world, channel scale guaranteed visibility. Bigger retailers led to higher search rank and better shelf placement. This meant a bigger audience and more sales.

In the AI world, trust signals matter more than the channel itself. A smaller brand with stronger trust signals can beat a larger brand with weaker ones.

Similar to the dominance of Reddit data shown above, we've found that each named badge, label, or third-party endorsement is a high-value mention. This is why trust signals matter more than ever for product discovery, and retailers and online marketplaces that give these product labels and endorsements are winning.


"Every time your product is endorsed, you increase your chances of being cited by AI."

The chart below shows the magnitude. Verified data on PDPs in ChatGPT were over 60 times more likely to be selected than PDPs with unverified data or fluffy marketing language and claims. Diving into the importance of trust signals even further, promotional tags like "scarcity" or "sponsored" not only failed to help products with selection and ranking, they often hurt.

15-66x more likely: Verified products were **66x more likely to be selected by ChatGPT**. Promotional labels underperformed verified labels by **10-58x**.

This verification bias persists across all models. In selection trials across 8 well-known products, the ones with verified data were chosen 28.8-32.3% of the time they were presented — or 2.3-2.6X random chance. Products with no verification were chosen only 0.5-1.7% of the time they were presented — giving products with verified data a 15-66X advantage.


Figure 2: Selection Rate by Label Type Across Models



All LLMs reward verification, with higher selection rates.

Our research also found that named verifiers matter. For example, in ChatGPT, Ulta Conscious Beauty, MadeSafe, EWG, and Target Clean all outperformed generic labels by 3-4.5x.

Figure 3: Named Verifier Performance vs Generic (ChatGPT)



Named verifiers outperformed generic signals by 3-4.5x.

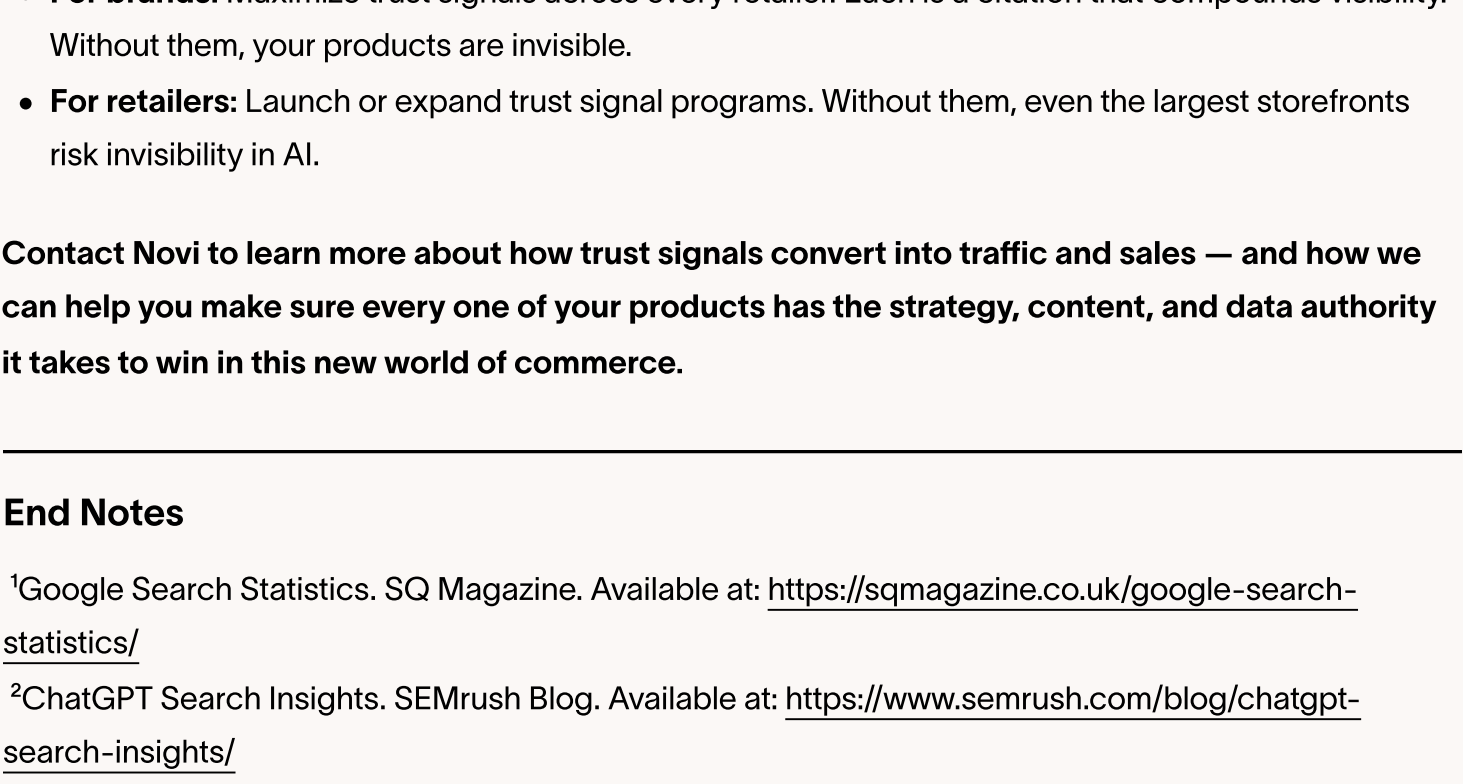
Perplexity reinforced our finding that Ulta badging is a top-performing verifier:

- Ulta's trust signal carried a **54.8% selection rate**, far above MadeSafe (36.1%), EWG (35.7%), or Target Clean (25.8%).
- Generic verifiers lagged at 16.5%.

In Claude and Gemini, Made Safe was the top-performing verifier (52.9% and 53.5% selection rate respectively), with EWG (49.1% and 42.9%) and Ulta (46.0% and 46.6%) just behind. Together, this suggests model-specific preferences for different sources of trust signals.

Our research also uncovered the effects of formatting. How you display your content is just as important as the content itself. For ChatGPT, pairing a badge or label with AI optimized text that reinforced the label had a 51.8% selection rate versus just 7.1% for products using optimized text alone. The pattern was the same with Perplexity and Gemini, with a 48% selection rate for badge + text for Perplexity and 43.4% for Gemini. This magnifies the need to provide credible proof and backing behind claims, not optimized marketing text alone.

Figure 4: Format Performance in Chat GPT



Badge + AI optimized text is **7.3x stronger than text only**.

Experimental confirmation

We've corroborated these results with independent research. The *What Is Your AI Agent Buying?* paper⁴ found that a single "Overall Pick" label doubled — and in some models, more than quadrupled — the odds of being chosen (from 10% to as high as 43%).

In that same study, they found seller-side description edits alone drove **+21.8% market share gains** for mousepads and **+16.3%** for office lamps.

It's clear that AI agents are not replicating human heuristics. They are amplifying credibility signals and ignoring the remaining fluff.

Strategic implications

For brands: Success requires volume of credibility. A single channel doesn't guarantee AI visibility; repeated trust signals across multiple channels do.

For retailers: It's clear that trust signals are the new currency of AI commerce. Without trust signals, your scale doesn't protect you. Your PDPs will be invisible to AI agents.

Brand Imperative: Maximize Trust Signal Coverage

AI has opened up the conversational aspects of search and consumers are asking questions that never existed in search before. They have moved beyond generic keyword searches to hyper-niche, personal and detailed questions. But with these specific and specialized searches come a new weight on relevance, quality and accuracy of answers.

Shoppers with hyper-personalized queries showed a +257% increase in selection when relevant verified data was present and optimized in product description content. Meanwhile in shorter, more typical SEO-like searches shoppers responded primarily to ratings. Your job is to make sure that every SKU has the strategy, content, and data authority to win at these questions.

257% lift: Shopper selection increased by 257% when relevant verified data was present and optimized in product descriptions.

Case Studies

Neutrogena vs. CeraVe

In a targeted case study focused on competing brands, our research found that third-party validation was the single strongest optimization lever and even more powerful than performance claims or ingredient lists.

When optimizing Neutrogena's PDP content for ChatGPT and infusing verified claims, its rate of selection grew +200%. On top of that, its closest competitor and the market leader, CeraVe, lost -5.5 points to Neutrogena despite incumbency and review dominance.

L'Oréal

In another targeted case study, we found that not all sources of credibility are created equal. Our experiments show that AI Agents weigh the authority of the source.

When L'Oréal's product pages carried third-party validation, selection odds rose sharply compared to baseline. When those same pages referenced a weaker or low-credibility source, selection odds fell.

Figure 5: L'Oréal Performance with Different Sources, ChatGPT

32% improvement with 3rd-party validation, -11% decline with weak/low-credibility source

Not only do trust signals matter, but the origination of those signals matters. Third-party validation consistently outperformed alternatives, proving that AI agents prefer claims backed by trusted, independent verifiers.

+32% lift: L'Oréal PDPs with Novi validation outperformed baseline by 32%. **Credibility compounds.** Winning in AI requires "signaling everywhere."

Retailer Imperative: Launch or Expand Trust Signal Programs

For retailers, the stakes are even higher. Both the opportunity, and risk, is enormous. Without trust signal programs, all of the products you sell risk being invisible in AI. With them, your site becomes the authoritative source AI agents cite.

Risks

- PDPs without trust signals are systematically down-weighted.
- Scale alone no longer protects your traffic.
- Poorly executed badges invite legal risk (FTC scrutiny, consumer trust erosion)

Opportunities

- **Authoritative source:** Embedding trust signals into PDP metadata ensures AI models recognize your credibility, cite your site and not competitors, and you capture disproportionate AI-driven traffic.
- **Sales leverage:** Retailer-issued endorsements like "Overall Pick" have been shown to deliver the equivalent of **65-138% price premium gains** in AI models (ACES study).
- **Human shoppers still care:** Amazon's Climate Pledge Friendly program delivered **12-15% lifts on average, up to 70% in some categories**.
- **Retailer advantage:** Trust signals aren't just for brands; they elevate the retailer as the trusted storefront in AI commerce.

In trials with ChatGPT, products with verification were selected 26.8X more often than those with generic "overall pick" labeling, and those with no trust signals were selected < 1% of the time. . In one lipstick catalog test, five of eight tested SKUs were never selected without trust signals.

<1%: ChatGPT's selection rate for products without verification was just 0.5%.

In perplexity trials, verification bias remained strong (selection rate was 15.9x higher for those with verified signals) but position 1 merchandising doubled visibility, highlighting the synergy between retailer placement and verification. And the ACES study corroborates this: the "Overall Pick" labels produced price-equivalent gains of 65-138% depending on the model.

Strategic Implications for Retailers

Trust signals are no longer optional. They are the gateway to AI visibility. Retailers hold the keys: the choices they make on product labeling, verification, and PDP metadata directly determine whether their assortments are surfaced by AI agents.

Key Takeaways

AI has rewritten the rules of commerce. Visibility is no longer about scale, rank, or placement. It's about trust.

- **For brands:** Maximize trust signals across every retailer. Each is a citation that compounds visibility. Without them, your products are invisible.
- **For retailers:** Launch or expand trust signal programs. Without them, even the largest storefronts risk invisibility in AI.

Contact Novi to learn more about how trust signals convert into traffic and sales — and how we can help you make sure every one of your products has the strategy, content, and data authority it takes to win in this new world of commerce.

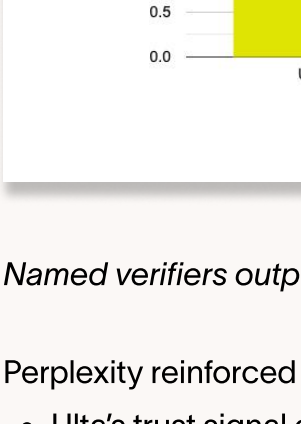
End Notes

¹Google Search Statistics. SQ Magazine. Available at: <https://sqmagazine.co.uk/google-search-statistics/>

²ChatGPT Search Insights. SEMrush Blog. Available at: <https://www.semrush.com/blog/chatgpt-search-insights/>

³Smith, E. The Ultimate Guide to AEO. Lenny's Newsletter. Available at: <https://www.lennysnewsletter.com/p/the-ultimate-guide-to-aeo-ethan-smith>

⁴"What is Your AI Agent Buying? Evaluation, Implications, and Emerging Questions for Agentic E-Commerce" (Allouah, Besbes, Figueroa, Kanoria, & Kumar, arXiv:2508.02630, 2025).



About the Author

Kimberly Shenk | Co-Founder and CEO | Novi Connect

Kimberly Shenk is the co-founder and CEO of Novi. She holds a bachelor's in mathematics and operations research from the US Air Force Academy, and a master's in operations research and data science from MIT. Before becoming a serial entrepreneur, she served as a senior data scientist in the US Air Force and then led data and product organizations at Eventbrite and Domino Data Labs. At Novi, she's applying her deep expertise in data science to build the trust infrastructure for commerce, helping brands turn verified product data into visibility, credibility, and growth so that both consumers and AI agents can make better decisions in the age of AI shopping.