

The Trust in AI Commerce Report

Edition 2 · October 2026

Product.ai's recurring survey of how U.S. shoppers use, trust and check AI product recommendations. Edition 2 surveyed 1,493 U.S. shoppers from August 27 to September 1, 2026 and repeats 12 questions from the June 2026 edition word for word.

PRODUCT.AI RESEARCH · CONTRIBUTORS: MICHAEL QUOC, DAKOTA NUNLEY, KAREN SPINNER · OCTOBER 14, 2026

PRODUCT.AI/RESEARCH/TRUST-IN-AI-COMMERCE-REPORT · CC BY 4.0

Edition 2 · Published October 14, 2026 · 1,493 U.S. shoppers · Survey fielded August 27 to September 1, 2026. This is the second edition of Product.ai's recurring survey of U.S. shoppers. The first edition, published in June 2026 with 1,463 U.S. online shoppers, is [available in full on its own page](#), and the questions the two editions share are compared throughout.

Key takeaways

- 1. Nine in ten shoppers who use AI for research check AI answers before buying. Product.ai's Edition 2 survey found that 89.9% of U.S. online shoppers who used an AI assistant for product research verified its recommendation through another source before buying, up from 85.7% in the June 2026 edition.** 41.0% always verify, 48.9% sometimes verify, and 10.1% trust the recommendation as given, down from 14.3% in the first edition. Checking AI answers finds errors. 89.9% of shoppers who verify have found AI mistakes at least once, and 60.3% of those changed course the last time it happened, abandoning the purchase, switching products, or dropping AI for that category. (*n=517 AI users among 1,247 online shoppers on the trend base; verifiers n=493 of 550 AI users, error-finders n=443; source: §2.2, §1.2, §2.4, §2.5*)
- 2. Four in ten U.S. online shoppers use AI to research purchases. In the 2026 Trust in AI Commerce Report, 41.5% used an AI assistant to research a purchase in the past 90 days, against 42.6% in the June 2026 edition.** The 1.1-point difference is smaller than the 4-point change threshold the research team set before fielding, so the report treats adoption as unchanged. (*n=1,247 online shoppers; source: §1.1*)
- 3. One in five shoppers who use AI for research has let an assistant buy something for them. In Product.ai's Edition 2 survey, 21.6% of AI users report letting an AI assistant complete a purchase start to finish; a further 27.8% reached checkout with an assistant and paid themselves.** Completion is 28.8% among shoppers who always verify AI recommendations and 8.8% among those who never do. (*n=550 AI users; source: §5.1, §5.3*)
- 4. Shoppers research with Gemini and ChatGPT, and most use more than one assistant. In Product.ai's survey, Google Gemini (63.5% of AI users) and ChatGPT (60.0%) lead research use, and the two sit within two points of each other on the most-trusted ranking (43.5% and 42.2%). 64.0% of AI users**

rely on two or more assistants, and the heaviest users delegate most. 31.5% of those using three or more have completed an AI purchase, against 13.7% of single-assistant users. (*n=550 AI users; source: §3.1, §3.2, §5.4*)

5. **Most AI users will not let an agent spend on its own. Product.ai found that 37.5% of AI users would let an AI agent spend nothing at all without checking with them first, and 43.1% would not give an agent access to a payment method.** “Only if it asks first” is the most common answer for every one of the seven agent tasks the survey put to shoppers. (*n=550 AI users; source: §4.2, §4.3*)
6. **AI is already part of holiday shopping. In the 2026 Trust in AI Commerce Report, 36.5% of AI users had used an assistant for holiday shopping by the time the survey fielded in late August, and 71.5% expect to this season, 27.6% definitely and 43.8% probably. Among shoppers who have let an assistant complete a purchase, 83.2% had already used one for holiday shopping.** (*n=550 AI users; completers n=119; source: §5.5*)
7. **Hiding a sponsorship costs more trust than disclosing one. Product.ai’s survey found that 50.4% of AI users would trust an assistant less after discovering a paid recommendation they were not told about, against 39.8% for a recommendation clearly labeled as sponsored. The 10.5–point difference is the Disclosure Premium.** 26.9% say a clear label would make them trust the assistant more. (*n=550 AI users; source: §8.2*)
8. **Shoppers trust AI least of seven product–research sources. AI assistants score 4.12 out of 7 in Product.ai’s October 2026 edition, the same score as in the June 2026 edition.** In this edition, Reddit and forums scored 4.18 and YouTube creators 4.30. Friends and family lead both editions. (*n=1,247 online shoppers; Edition 1 n=1,462; source: §7.1, §1.1*)
9. **“Verified” has no shared meaning. Asked by Product.ai what the word means when a shopping site or AI assistant uses it, 14.0% of shoppers described an independent check or test, 14.9% said they did not know, and the largest group offered synonyms for “real” or “legit.”** (*n=1,493, open text; source: §9.1*)

How this survey was done

Product.ai surveyed 1,493 U.S. adults who shop for products, online or in stores, from August 27 to September 1, 2026, on a general–population panel. Twelve questions repeat the June 2026 edition word for word, and comparisons use the 1,247 respondents who shop online, the group the first edition sampled. Results are unweighted, with panel quotas on age, gender, and region. [Full methodology.](#)

Executive summary

The Trust in AI Commerce Report is Product.ai’s recurring measurement of how U.S. shoppers use, trust, and check AI product recommendations. Edition 2 was fielded from August 27 to September 1, 2026 on a general–population panel and produced 1,493 complete responses. Product.ai first published this survey in June 2026, with 1,463 U.S. online shoppers. Edition 2 asks 12 of those questions again, word for word, so each change between the two editions compares like with like. Five new topics were added: which AI assistants shoppers trust, sponsored content inside AI answers, the permissions shoppers would give a shopping agent, completed AI purchases, and holiday shopping.

Most measures repeated from the first edition did not change. 41.5% of U.S. online shoppers used AI to research a purchase in the past 90 days, against 42.6% in the first edition, a difference smaller than the 4-point change threshold set before fielding. 42.2% would not spend more than \$25 on an AI-recommended product without checking another source, against 41.8% in the first edition. AI remained last of seven product-research sources, at 4.12 out of 7 in both editions. Confidence in shoppers' own most recent purchase, retailer trust, and the leading purchase category did not change. Verification was the only measure that changed, and it rose. 89.9% of AI users now verify an AI recommendation before buying, up from 85.7%.

Against that stable background, 21.6% of AI users report having let an AI assistant complete a purchase from start to finish, and 66.9% have at least started one. The shoppers doing this are the ones who check most. Completion is 28.8% among shoppers who always verify, 18.5% among those who sometimes verify, and 8.8% among those who never verify. At the same time, 37.5% of AI users would let an agent spend nothing unsupervised and 43.1% refuse it a payment method. Shoppers who supervise AI are the ones delegating to it, and most limit or refuse agents' access to their money.

Google Gemini (63.5% of AI users) and ChatGPT (60.0%) lead research use and sit within two points of each other on the trust ranking. 64.0% of AI users rely on two or more assistants. Concealed sponsorship costs 10.5 points more trust than labeled sponsorship. Three-quarters of AI users already knew or suspected that an AI answer can repeat a brand's marketing as fact, and half still agree that AI shopping answers are unbiased. Asked what "verified" means, one shopper in seven described an independent check.

The first edition included only shoppers who buy online. Edition 2 also kept 246 respondents who shop only in physical stores. They use AI for shopping research at 13.4% and give retailers a Net Retailer Trust Score of -52.8, against -23.7 among online shoppers. The score is the share who rate retailers 9 or 10 out of 10 for having their best interest in mind, minus the share who rate them 0 to 6. They enter this report as a separate segment and never inside a trend number.

Findings are stated with their denominators throughout. Interpretation is confined to §11, "Reading the numbers."

How this edition compares with the June 2026 edition


Product.ai published the first edition of this survey in June 2026. Edition 2 asked 12 of its questions again, with identical wording and the same kind of respondent, U.S. adults who shop online. Comparisons use the 1,247 Edition 2 respondents who shop online, the same group the first edition sampled, so that keeping [store-only shoppers](#) in Edition 2 does not distort the trend. The report calls this group the trend base. Before fielding, the research team set a rule for what counts as change. A move of about 4 points or more after design effects is reported as a change. Smaller moves are reported as held steady.

1.1 What changed between the two editions

MEASURE	JUNE 2026 EDITION	OCTOBER 2026 EDITION	CHANGE	READING
Used AI to research a purchase, past 90 days (Q12)	42.6% (n=1,463)	41.5% (n=1,247)	-1.1	Held steady
Verified the AI recommendation before buying, of AI users (Q13)	85.7% (n=623)	89.9% (n=517)	+4.2	Moved, at the threshold
Always verify, of AI users	44.5%	41.0%	-3.5	Held steady
Never verify, of AI users	14.3%	10.1%	-4.2	Moved, at the threshold
Would spend under \$25 on an AI recommendation without checking (Q23)	41.8%	42.2%	+0.4	Held steady
Used an AI assistant to research the most recent \$50+ purchase (Q7)	20.3%	17.5%	-2.8	Held steady
Last-purchase confidence, 0-10 (Q5)	6.93 (n=1,453)	7.03 (n=1,243)	+0.10	Held steady
Net Retailer Trust Score (Q10)	-25	-23.7	+1.3	Held steady
Retailer trust mean, 0-10 (Q10)	6.35	6.40	+0.05	Held steady
AI tool or assistant, mean trust (Q11, 1-7)	4.12	4.12	0.00	Held steady
Seven-source trust, spread top to bottom	1.15	1.22	+0.07	Held steady
Per-category research confidence, 18-category mean (Q9, 1-7)	4.88	4.95	+0.07	Held steady
Trust in AI recommendations, 18-category mean (Q22, 1-7)	4.01	4.01	0.00	Held steady
Lowest AI-trust category	Baby products, 3.82	Baby products, 3.79	-0.03	Held steady
Most recent \$50+ purchase category (Q4)	Apparel, 30.4%	Apparel, 30.1%	-0.3	Held steady

Verification rose from 85.7% in June to 89.9% in October

Share of AI users who checked the recommendation through another source before buying.




N=623 EDITION 1 · N=517 EDITION 2, TREND BASE · FIELDED AUG 27 TO SEP 1, 2026 · PRODUCT.AI RESEARCH

The share of AI users who checked an AI recommendation through another source before buying, the June 2026 edition (n=623) and the October 2026 edition (n=517). This chart gains a point with each edition.

AI research adoption held at about 42% across both editions

Share of U.S. online shoppers who used an AI assistant to research a purchase in the past 90 days.




N=1,463 EDITION 1 · N=1,247 EDITION 2, TREND BASE · FIELDED AUG 27 TO SEP 1, 2026 · PRODUCT.AI RESEARCH

The share of U.S. online shoppers who used an AI assistant to research a purchase in the past 90 days, the June 2026 edition (n=1,463) and the October 2026 edition (n=1,247).

The \$25 threshold held: 41.8% in June, 42.2% in October

Share who would spend under \$25 on an AI-recommended product without checking another source.



N=1,463 EDITION 1 · N=1,247 EDITION 2, TREND BASE · FIELDLED AUG 27 TO SEP 1, 2026 · PRODUCT.AI RESEARCH

The share of online shoppers who would spend under \$25 on an AI-recommended product without checking another source, both editions.

1.2 More AI users now check the answer before buying

Verification moved 4.2 points, from 85.7% to 89.9% of AI users, and the share who never verify fell from 14.3% to 10.1%. Both changes are at the pre-registered threshold. The share of AI users who act on a recommendation without checking fell from one in seven to one in ten.

1.3 Adoption, spending limits, and retailer trust held steady

13 of the 15 measures in the §1.1 table held steady, including AI adoption, the share who would not spend more than \$25 unchecked, purchase confidence, retailer trust, the leading purchase category, AI's trust score, and both 18-category grids. These measures show neither a market warming to AI nor one turning away from it. Questions asked for the first time in this edition have no earlier benchmark.

(Trend base: online shoppers, n=1,247 in Edition 2; Edition 1 n=1,463. Change threshold pre-registered before fielding; source: Methodology, Change threshold.)

The Verification Reflex: How shoppers use AI and check its answers

The Verification Reflex is the report's name for a shopper checking an AI recommendation through another source before acting on it. Edition 2 measured whether shoppers used AI for research, whether they verified, where they checked first, how often checking found the AI wrong, and what they did next.

2.1 Four in ten online shoppers used AI to research a purchase

41.5% of U.S. online shoppers used an AI assistant to research a product they were considering in the past 90 days (517 of 1,247). On the full sample, which includes store-only shoppers, the figure is 36.8% (550 of 1,493). The AI-user chapters that follow use the 550.

The age pattern of adoption changed between editions. Under-35 shoppers held near their first-edition levels while the 35-44 band fell and the two oldest bands rose:

AGE	JUNE 2026 EDITION	OCTOBER 2026 EDITION (TREND BASE)
18-24	55.6% (n=171)	54.8% (n=146)
25-34	57.5% (n=259)	52.4% (n=212)
35-44	58.2% (n=239)	46.4% (n=237)
45-54	44.0% (n=232)	38.0% (n=192)
55-64	29.6% (n=247)	38.0% (n=200)
65 or older	20.6% (n=315)	25.8% (n=260)

In the June 2026 edition, adoption dropped 14.2 points between the 35-44 and 45-54 age bands. In this edition the drop is 8.4 points. The age gradient is still significant (chi-square 78.5, p under .001) and it is flatter than it was.

(Adoption: n=1,247 online shoppers, 41.5%; full sample n=1,493, 36.8%. Age cells 146 to 260; source: §2.1, §10.3.)

2.2 Nine in ten AI users checked the recommendation before buying

Among AI users on the trend base, 89.9% verified the assistant's recommendation through another source before making a purchase. 41.0% always verify and 48.9% sometimes verify; 10.1% do not. On the full AI-user base of 550 the figures are 89.6%, 40.4%, 49.3%, and 10.4%.

The 89.9% describes the 517 people who used AI for product research. It does not describe all 1,247 online shoppers, most of whom received no AI recommendation to verify.


(Verification: n=517 AI users on the trend base; 41.0% always + 48.9% sometimes; Edition 1 85.7% of 623; source: §2.2.)

2.3 Most shoppers who verify check customer reviews first

This edition is the first to ask where shoppers go to check an AI answer. The 493 AI users who verify (called verifiers in the tables) were asked where they looked first the last time they double-checked an AI recommendation:

FIRST PLACE CHECKED	SHARE OF VERIFIERS (N=493)
Online customer reviews	36.7%
A search engine	21.9%
The brand's or retailer's own website	14.0%
Friends or family	13.4%
Another AI assistant	5.7%
YouTube or video reviews	4.7%
An expert publication	1.8%
Somewhere else	1.0%

Reviews and search together account for 58.6% of first checks. Checking one assistant against another is a real behavior and a small one, at 5.7%. Shoppers verify AI against customer reviews while rating their belief that reviews reflect product quality at 4.56 out of 7 (§7.3).



Where verifiers looked first the last time they double-checked an AI recommendation, n=493 AI users who verify.

(First verification stop: n=493 verifiers among 550 AI users; source: §2.3.)

2.4 Most shoppers who verify have found AI mistakes

Asked how often checking has shown the AI's answer to be wrong or out of date, shoppers who verify said:

FREQUENCY	SHARE OF VERIFIERS (N=493)
Never	10.1%
Rarely	28.4%
Sometimes	39.4%
Often	14.6%
Every time	7.5%

89.9% of shoppers who verify have caught the AI wrong at least once. 61.5% catch it at least sometimes, and 22.1% often or every time.

(Error frequency: n=493 verifiers; source: §2.4.)

2.5 After an error, most shoppers changed what they bought or stopped using AI


The 443 shoppers who verify and had found an AI answer wrong were asked what they did next, the last time it happened:

ACTION AFTER FINDING THE AI WRONG	SHARE (N=443)
Decided not to buy at all	26.2%
Bought a different product	22.3%
Did nothing in particular	19.2%
Bought it anyway	17.2%
Stopped using AI for that kind of shopping	11.7%
Reported not double-checking	3.4%

60.3% changed course. They abandoned the purchase, switched products, or stopped using AI for that category. 36.3% carried on, either buying anyway or doing nothing in particular.

Six in ten shoppers changed course after catching an AI answer wrong

What AI users did next the last time checking showed the assistant was wrong.



N=443 AI USERS WHO FOUND AN ERROR · FIELDLED AUG 27 TO SEP 1, 2026 · PRODUCT.AI RESEARCH

What AI users did next the last time checking showed the assistant was wrong, n=443.

(Behavior after error: n=443 error-finders among 493 verifiers; source: §2.5.)

2.6 Most AI users believe the assistant found the best price

Two beliefs about the most recent AI-assisted purchase were measured separately. 68.2% of AI users believe the assistant found the best available price. 56.0% believe it saved them money. 22.0% said it did not save them money and 22.0% were not sure. Among shoppers who believe AI saved them money, 90.3% also believe it found the best price.

(Savings beliefs: n=550 AI users; source: §2.6.)

2.7 Summary of findings

- 41.5% of online shoppers used AI to research a purchase in the past 90 days (§2.1).
- 89.9% of AI users verified the recommendation before buying, up from 85.7% (§2.2).
- 36.7% of shoppers who verify check customer reviews first; 5.7% check another assistant (§2.3).
- 89.9% of shoppers who verify have found the AI wrong at least once; 60.3% of those changed course by abandoning the purchase, switching products, or dropping AI for that category (§2.4, §2.5).
- 68.2% believe AI found the best price; 56.0% believe it saved them money (§2.6).

Which assistants shoppers use and trust


Edition 2 asked the 550 shoppers who used AI for product research, for the first time, which assistants they used, which they trust most and least for shopping, and why.

3.1 Gemini and ChatGPT lead, and most shoppers use more than one assistant

ASSISTANT USED IN THE PAST 90 DAYS	SHARE OF AI USERS (N=550)
Google Gemini	63.5%
ChatGPT	60.0%
Amazon's shopping assistant	27.5%
Walmart's shopping assistant	22.5%
Microsoft Copilot	22.4%
Claude	17.6%
Perplexity	11.3%

Respondents could name more than one. 64.0% of AI users rely on two or more assistants for product research, and the average is 2.25. Shoppers who use three or more assistants have completed an AI purchase at 31.5%, against 13.7% among those who use one (n=200 and n=183; p under .001; completed purchases are covered in §5). Grok and Meta AI appeared repeatedly as write-ins and will be listed in Edition 3.

Claude and Perplexity are almost never a shopper's only assistant. Of 550 AI users, six use Claude alone and four use Perplexity alone, against 77 for Gemini and 66 for ChatGPT. Claude users use 3.79 assistants on average and Perplexity users 3.84.



Share of AI users by number of assistants used, with the completed-purchase rate for each group, n=550.

(Assistants used: n=550 AI users, multi-select; stack analysis n=183 single, n=200 three or more; source: §3.1.)


3.2 Shoppers trust Gemini and ChatGPT most and Perplexity least

AI users marked which of seven assistants they trust most and which they trust least for shopping recommendations. A shopper could name more than one assistant in each column, so each column sums past 100%. The net score is the trust-most share minus the trust-least share.

ASSISTANT	TRUST MOST	TRUST LEAST	NOT CHOSEN EITHER WAY	NET
Google Gemini	43.5%	15.3%	41.1%	+0.28
ChatGPT	42.2%	19.1%	38.7%	+0.23
Amazon's shopping assistant	29.7%	20.9%	49.4%	+0.09
Microsoft Copilot	17.8%	23.8%	58.5%	-0.06
Walmart's shopping assistant	22.2%	29.6%	48.2%	-0.07
Claude	11.0%	31.3%	57.7%	-0.20
Perplexity	8.2%	35.2%	56.6%	-0.27

Three assistants have a positive net score. Walmart's assistant is net negative despite research usage close to Amazon's. Claude and Perplexity draw the highest trust-least shares and the lowest trust-most shares.

The ranking should be read next to §3.1. The order of trust matches the order of usage for six of the seven assistants, and between 38.7% and 58.5% of shoppers did not pick each assistant either way, because shoppers concentrate their most- and least-trusted choices on the assistants they use. An assistant a shopper has not used tends to end up in the least-trusted column. The ranking measures familiarity as much as trust, and §11 returns to it.



Share of AI users naming each assistant among those they trust most for shopping, more than one allowed, n=550.

(Most and least trusted: n=550 AI users; values from the survey platform’s standard report, see Methodology, Data sources; source: §3.2.)

3.3 Accuracy is the most common reason to trust or mistrust an assistant

Asked in their own words what makes them trust or mistrust an AI assistant for shopping recommendations, 17.1% of the 549 who answered spoke about accuracy, 9.3% about whether the assistant agreed with their own research or with reviews, 6.7% about bias, advertising, or commercial motive, and 5.8% about reliability. 4.0% said an assistant lacks human judgment about fit, 3.8% praised speed and breadth, and 2.0% expressed general distrust of AI. About 54% of answers were too short or too vague to code.

(Open text: n=549 AI users; theme coding in Appendix B of the codebook; source: §3.3.)

The \$25 threshold and the Agent Allowance: How much shoppers let AI spend unchecked

The \$25 AI Autonomy Threshold, first measured in the June 2026 edition, is the most a shopper would spend on an AI-recommended product without checking another source. The Agent Allowance, new in Edition 2, is the most a shopper would let an AI agent spend on their behalf without asking.

4.1 42% of online shoppers won’t spend over \$25 on an unchecked AI pick

42.2% of U.S. online shoppers would not spend more than \$25 on a product recommended by an AI assistant without checking another source first (n=1,247). In the June 2026 edition the figure was 41.8%. On the full sample, including store-only shoppers, it is 46.1%.


MOST A SHOPPER WOULD SPEND ON AN AI RECOMMENDATION WITHOUT CHECKING	SHARE, FULL SAMPLE (N=1,493)
Under \$25	46.1%
\$25-\$50	20.0%
\$51-\$100	14.9%
\$101-\$250	9.1%
\$251-\$500	5.3%
Over \$500	4.6%

(Threshold: 42.2% on the trend base, n=1,247; Edition 1 41.8% of 1,463; full-sample distribution n=1,493; source: §4.1.)

4.2 Most AI users would let an agent spend little to nothing on its own

Asked the most an AI agent should be able to spend without checking with them, 37.5% of AI users said nothing at all. 23.5% capped it under \$25 and 27.1% at \$25-\$100. 10.0% would allow \$100-\$500 and 2.0% more than \$500. 61.0% cap the agent at under \$25 or at zero, a stricter line than the \$25 recommendation threshold.

Comfort with a fixed monthly amount an agent could spend is low. On a 7-point scale the mean is 3.30; 54.6% sit on the uncomfortable side and 28.2% on the comfortable side. Among shoppers who have completed an AI purchase the mean is 4.90; among those who have never started one it is 2.39.



The most an AI shopping agent should be able to spend without checking, n=550 AI users.

(Agent Allowance: n=550 AI users; monthly-allowance comfort mean 3.30 of 7; source: §4.2.)

4.3 Most AI users refuse or limit an agent’s access to payment

43.1% of AI users would not give an AI agent access to a payment method. 37.1% would with limits. 19.8% would without limits. The 441 who limit or refuse were asked why, and could pick more than one reason:

REASON FOR LIMITING OR REFUSING PAYMENT ACCESS	SHARE (N=441)
I want control over what it buys	60.3%
Fraud or security worries	56.2%
It might make expensive mistakes	38.8%
I do not trust the company that built it	24.0%
I do not trust the shops it would buy from	23.6%

Control outranks fraud. If an agent bought the wrong thing with their money, 45.8% of AI users would hold themselves most at fault and 44.7% the company that built the agent. The shop (5.6%) and the payment provider (3.8%) barely register.


(Payment access: n=550 AI users; reasons n=441; blame n=550; source: §4.3.)

4.4 Shoppers want the agent to ask first, whatever the task

For seven tasks a shopping agent could perform, AI users were asked whether they would let it act without asking, only if it asks first, or never:

TASK	WITHOUT ASKING	ONLY IF IT ASKS FIRST	NEVER
Apply a promo code	38.9%	44.5%	16.5%
Track prices	35.6%	49.8%	14.5%
Research products	34.5%	47.8%	17.6%
Return an item	20.0%	51.6%	28.4%
Reorder something bought regularly	19.6%	46.5%	33.8%
Add things to a cart	18.7%	47.1%	34.2%
Place an order	16.9%	44.2%	38.9%

“Only if it asks first” is the most common answer on every row. Shoppers will let an agent apply a code or track a price far more readily than place an order.



Share of AI users who would let an agent act without asking, only if it asks first, or never, across seven tasks, n=550.

(Permission map: n=550 AI users; source: §4.4.)


Who has let an AI assistant complete a purchase

Edition 2 is the first edition to ask whether a shopper has let an AI assistant complete a purchase. The question read: “Have you ever let an AI assistant complete a purchase for you, start to finish?” Respondents chose between yes, reaching checkout but paying themselves, starting but stopping before paying, and never. A “yes” means the shopper reports the assistant completing the purchase through payment; it may include a checkout the assistant opened and the shopper confirmed inside it. The report uses “completed purchase” in that self-reported sense throughout.

5.1 One in five AI users has let an assistant complete a purchase

HAS LET AN AI ASSISTANT COMPLETE A PURCHASE	SHARE OF AI USERS (N=550)
Yes, start to finish	21.6%
Reached checkout, but paid myself	27.8%
Started one, stopped before paying	17.5%
Never	33.1%

21.6% of AI users have completed an AI purchase, which is 8.0% of all 1,493 shoppers surveyed. 66.9% have at least started an agent-led checkout. The largest drop is at payment. 45.3% of AI users reached or started a checkout and did not let the assistant pay.



The completed-purchase funnel among AI users, n=550.

(Completed purchases: n=550 AI users, 119 completers; source: §5.1.)

5.2 More shoppers have let an agent buy than say they would allow it

16.9% of AI users say they would let an agent place an order without asking (§4.4). 21.6% report already having let one complete a purchase. Behavior runs ahead of the strictest stated permission.


(Stated versus actual: n=550 AI users; source: §4.4, §5.1.)

5.3 Shoppers who always verify are the most likely to let an agent buy

Completion rises with verification habit. Among AI users who always verify, 28.8% have completed an AI purchase. Among those who sometimes verify, 18.5%. Among those who never verify, 8.8%.

VERIFICATION HABIT	COMPLETED AN AI PURCHASE
Always verify (n=222)	28.8%
Sometimes verify (n=271)	18.5%
Never verify (n=57)	8.8%

The 119 shoppers who have completed an AI purchase differ from AI users at large on every permission measure. 42.9% of them would let an agent place an order without asking, against 16.9% of all AI users. 54.6% would give an agent unlimited payment access, against 19.8%. 11.8% would allow an agent to spend nothing, against 37.5%. They also have a lower \$25 threshold (9.2% cap at under \$25, against 23.9% of AI users who have never started an AI purchase) and a higher monthly-allowance comfort (4.90 against 2.39).



Completed AI purchases by verification habit among AI users: always 28.8%, sometimes 18.5%, never 8.8%.


(Completion by verification: n=222, 271, 57; completer profile n=119; source: §5.3.)

5.4 ChatGPT and Gemini completed most of the purchases

Shoppers who have completed an AI purchase were asked which assistant they let complete it and could name more than one:

ASSISTANT	SHARE OF COMPLETERS (N=119)
ChatGPT	60.5%
Google Gemini	47.9%
Claude	16.0%
Amazon's shopping assistant	14.3%
Walmart's shopping assistant	14.3%
Microsoft Copilot	11.8%
Perplexity	5.9%

The base is 119 people, so shares for the smaller assistants are directional. ChatGPT's lead over Gemini in this group (60.5% against 47.9%) is not statistically significant on this base; 43 people named ChatGPT and not Gemini, 28 named Gemini and not ChatGPT, and 29 named both. The report does not claim that one frontier assistant closes purchases better than the other.



Which assistant completers let complete a purchase, multi-select, n=119 completers.

(Completing assistant: n=119 completers, multi-select; source: §5.4.)

5.5 More than a third of AI users have used an assistant for holiday shopping

36.5% of AI users had already used an AI assistant for holiday shopping when the survey fielded in late August 2026; 5.5% said they do not do holiday shopping. 71.5% expect to use one this season: 27.6% definitely and 43.8% probably. Among shoppers who have completed an AI purchase, 83.2% had already used an assistant for holiday shopping.

(Holiday: n=550 AI users; completers n=119; source: §5.5.)


Purchase confidence (PRCI): How sure shoppers are of their own research

The Purchase Research Confidence Index has two parts. The first is a single 0–10 rating of how confident a shopper was that they found the best product before completing their most recent online purchase of \$50 or more. The second is a set of 18 ratings, one per product category, of how confident the shopper feels finding the best product when shopping in that category.

6.1 Confidence in the last purchase held at 7 out of 10

The mean confidence rating for the most recent \$50+ online purchase is 7.03 out of 10 on the trend base (n=1,243), against 6.93 in the June 2026 edition, a difference reported as held steady. On the full sample the mean is 6.88; among store-only shoppers it is 6.12. 23.7% of respondents rated their confidence a 10; 49.1% rated it 7 or below.

The most recent \$50+ purchase was in Apparel for 30.1% of online shoppers and Electronics for 17.1%, against 30.4% and 19.8% in the first edition.



Mean confidence in the most recent \$50+ online purchase, 0–10, online shoppers, both editions.

(Last-purchase confidence: n=1,243 online shoppers with a rating; Edition 1 n=1,453; source: §6.1.)

6.2 Shoppers feel most confident buying apparel and least confident buying baby products

On the trend base the cross-category mean is 4.95 out of 7. Apparel leads at 5.38, with Electronics (5.24) and Grocery and food delivery (5.22) behind it. Baby products and child safety is the floor at 4.51, 0.44 points below the mean. On the full sample the mean is 4.82 and the range runs from 4.38 (Baby products) to 5.21 (Apparel).

CATEGORY (FULL SAMPLE, 1-7)	MEAN	CATEGORY	MEAN
Apparel	5.21	Furniture	4.78
Electronics	5.09	Mattress and bedding	4.77
Grocery and food delivery	5.07	Home improvement tools	4.72
Skincare	5.03	Video games	4.69
Kitchen appliances	5.03	Automotive accessories	4.68
Fashion and clothing	4.99	Beauty	4.60
Supplements and health	4.92	Outdoor and camping gear	4.56
Pet products	4.89	Fitness equipment	4.46
Toys and games	4.83	Baby products and child safety	4.38

The June 2026 edition found the same order, with Apparel highest and Baby products lowest. The 18-category mean was 4.88 in the first edition and 4.95 on the Edition 2 trend base, a difference reported as held steady (§1.1).

(Per-category confidence: n=1,493 full sample; trend base n=1,247; source: §6.2.)

6.3 Baby products rank last for both confidence and AI trust

Baby products and child safety ranks last on the confidence index in both editions, and last on the AI-trust grid as well (§10.1). The category with the highest stakes is the one where shoppers feel least equipped and least willing to trust AI. Who buys in that category, and how they use AI, is reported in §10.2.

Which information sources shoppers trust


Shoppers rated how much they trust seven sources of information when researching a product they might buy online. Edition 2 reports the grid on its native 1-7 scale.

7.1 AI ranks last of seven information sources in both editions

INFORMATION SOURCE	MEAN TRUST, 1-7 (TREND BASE)	RANK, EDITION 2	EDITION 1 MEAN (RANK)
Friends or family	5.34	1	5.27 (1)
Online customer reviews	4.91	2	4.96 (2)
Brand or retailer website	4.81	3	4.83 (4)
Expert publications or professionals	4.80	4	4.89 (3)
YouTube creators or reviewers	4.30	5	4.30 (5)
Reddit or online forums	4.18	6	4.22 (6)
AI tool or assistant	4.12	7	4.12 (7)

Friends and family lead by a clear margin in both editions. Brand websites and expert publications are separated by 0.01 and should be treated as tied for third. AI assistants rank last in both editions, at 4.12.

The sources fall into two tiers. A human and institutional tier (friends and family through expert publications) runs from 4.80 to 5.34. A digital-content tier (YouTube, Reddit, AI) runs from 4.12 to 4.30, a full half-point lower.




Mean trust in seven product-research sources, online shoppers, n=1,247.

(Trust hierarchy: n=1,247 online shoppers; full-sample means run 3.98 to 5.28 in the same order; source: §71.)

7.2 Trust in online retailers held steady

On the statement “Online retailers always have my best interest as a customer in mind,” the mean is 6.40 out of 10 on the trend base and the Net Retailer Trust Score is –23.7, against 6.35 and –25 in the June 2026 edition. The score subtracts the share rating 0 to 6 from the share rating 9 or 10. Among store-only shoppers it is –52.8 (§10.4).



Net Retailer Trust Score among online shoppers, both editions, with the store-only segment's –52.8 shown separately.

(Retailer trust: n=1,247 online shoppers; full sample mean 6.19, score –28.5; source: §7.2.)

7.3 Shoppers check reviews first but only moderately trust them

Asked how confident they are that online product reviews accurately reflect product quality, shoppers scored 4.56 out of 7 on the trend base (4.42 on the full sample). 24.3% chose the top two points and 12.9% the bottom two. Reviews rank second in comparative trust (§7.1) and are the first place shoppers who verify check an AI answer (§2.3), while drawing only moderate belief in their accuracy. The question asked about “online product reviews” without naming who wrote them; the trust hierarchy in §7.1 names “online customer reviews.” The report treats both as shopper-written reviews, not expert reviews, although the wording left that distinction to the respondent.

(Review accuracy belief: n=1,247; source: §7.3.)

The Disclosure Premium: How sponsored answers change trust

As sponsored recommendations begin to appear inside assistants, this edition added questions on whether shoppers have noticed them, how disclosure changes trust, how accurate shoppers believe AI answers are on three elements, and whether shoppers know that an AI answer can repeat a brand's marketing as fact.

8.1 Four in ten AI users have noticed a sponsored recommendation

41.5% of AI users have noticed a sponsored or paid recommendation inside an AI assistant's answer. 38.7% have not, and 19.8% are not sure.

(Sponsored recommendations noticed: n=550 AI users; source: §8.1.)

8.2 A hidden sponsorship costs more trust than a disclosed one

Two questions asked how trust would change if an assistant clearly labeled a recommendation as sponsored, and if the shopper found out afterwards that a recommendation had been paid for without being told.


CHANGE IN TRUST	CLEARLY LABELED	DISCOVERED AFTERWARDS
Trust it much less	16.5%	26.0%
Trust it somewhat less	23.3%	24.4%
No change	33.3%	30.2%
Trust it more	19.6%	12.9%
Trust it much more	7.3%	6.5%

39.8% would trust the assistant less over a labeled sponsorship. 50.4% would trust it less over a hidden one. The 10.5-point difference is the Disclosure Premium, the report's name for the trust an assistant keeps by telling shoppers a recommendation is paid. The share who would trust the assistant much less rises from 16.5% to 26.0% when the sponsorship is hidden. 26.9% of AI users say a clear label would raise their trust.

The premium is not uniform by age. Every band aged 55 and over penalizes a hidden sponsorship more than 20 points harder than a labeled one (22.9 points among 55–64s and 20.5 among those 65 and older). Shoppers aged 25–34 go the other way by 5.2 points. They mind the label slightly more than the concealment.

Product.ai's Edition 2 finds a hidden sponsorship costs 10.5 more points of trust than a labelled one

Extra trust lost when a sponsorship is hidden rather than labelled, in points.



N=550 AI USERS · FIELDLED AUG 27 TO SEP 1, 2026 · PRODUCT.AI RESEARCH

Share of AI users who would trust an assistant less over a labeled and a hidden sponsored recommendation, n=550, with the gap by age band.

(Disclosure Premium: n=550 AI users; age bands 73 to 118 AI users each, directional; source: §8.2.)

8.3 Shoppers believe AI most on specifications and least on price

AI users rated how accurate AI shopping answers usually are on three elements:

ELEMENT	MEAN, 1-7	BELIEVE ACCURATE (TOP THREE POINTS)
Product specifications and features	4.77	61.1%
Whether an item is in stock or still sold	4.61	55.8%
Price	4.45	53.5%

Product.ai's companion study, the [AI Shopping Divergence Study](#), measured how often four assistants, queried through their developer APIs, gave conflicting answers to the same shopping questions. Specifications are the element shoppers believe most, yet even plain spec and fact lookups, the question type on which the assistants agreed most, produced a confirmed conflict 75% of the time. The two studies are reported on their own pages; the paired reading is in §11.

(Believed accuracy: n=550 AI users; source: §8.3.)

8.4 Most AI users know AI can repeat marketing as fact, and half still call it unbiased

Asked whether, before the survey, they knew that an AI answer can repeat a brand's own marketing claims as if they were facts, 36.4% of AI users said they knew, 38.7% said they suspected, and 24.9% said they did not know. In the same survey, 48.7% agreed that "AI shopping answers are unbiased, not influenced by advertisers or brands" (mean 4.46 out of 7; 25.8% disagreed; 25.5% neutral).

Three-quarters of AI users already knew or suspected the problem. Half consider the assistant unbiased at the same time.

(Awareness that AI can repeat marketing as fact, and belief that AI is unbiased: n=550 AI users; source: §8.4.)

What “verified” means to shoppers

Every respondent was asked a new open question, in their own words: when a shopping site or an AI assistant says a product claim is “verified,” what do you think that actually means?

9.1 One shopper in seven thinks “verified” means an independent check


WHAT “VERIFIED” MEANS	SHARE OF SHOPPERS (N=1,493)
Do not know, or no clear meaning	14.9%
Someone independently checked or tested it	14.0%
It is authentic, legit, not a scam	9.9%
The claim is true or accurate	8.2%
It is safe, good quality, works as promised	7.8%
Real customers wrote the reviews	6.2%
Distrusts the label, or AI in general	5.7%
It was asserted by the seller, or is marketing	1.7%
Vague, tautological, or off-topic	remainder

Answers were coded by keyword frame and could match more than one theme. 54.6% of responses matched at least one substantive theme. The remainder were tautological (“it means it is verified”), vague, or unrelated.

One shopper in seven gave the word its usual meaning, that someone independent checked. About the same share said plainly that they do not know. The largest group used synonyms for legitimacy. Roughly 7% treated the label with suspicion, as a seller’s own claim or as marketing.

One shopper in seven thinks “verified” means someone checked

What shoppers said “verified” means on a shopping site or in an AI answer, coded open text.



N=1,493 SHOPPERS · FIELDLED AUG 27 TO SEP 1, 2026 · PRODUCT.AI RESEARCH

What U.S. shoppers said the word “verified” means on a shopping site or in an AI answer, coded open text, n=1,493.

(“Verified” open text: n=1,493, multi-assignment coding; representative verbatims in the codebook; source: §9.1.)

Differences by category, age, gender, income, and region

10.1 One in six shoppers refuses AI in every category

Every respondent rated how much they trust AI-generated product recommendations in each of 18 categories on a 1–7 scale. On the full sample the cross-category mean is 3.87, with Kitchen appliances at the ceiling (3.98) and Baby products and child safety at the floor (3.67). On the trend base the mean is 4.01, Kitchen appliances 4.13, Baby products 3.79.

CATEGORY (FULL SAMPLE, 1-7)	MEAN	CATEGORY	MEAN
Kitchen appliances	3.98	Video games	3.87
Electronics	3.97	Grocery and food delivery	3.87
Furniture	3.97	Automotive accessories	3.87
Pet products	3.94	Fitness equipment	3.83
Home improvement tools	3.93	Skincare	3.82
Fashion and clothing	3.92	Apparel	3.81
Mattress and bedding	3.90	Supplements and health	3.81
Outdoor and camping gear	3.90	Beauty	3.73
Toys and games	3.90	Baby products and child safety	3.67

The averages above combine two very different groups. 239 respondents, 16.0% of the sample, answered “do not trust at all” to every one of the 18 categories, and 95.4% of them had not used AI to research a purchase in the past 90 days. Among the other 1,254 respondents, every category is above the scale midpoint, from Baby products at 4.18 to Kitchen appliances at 4.55, in the same order. The rank order is a finding. The below-midpoint level of the averages is a property of combining refusers with everyone else, and the report does not claim it.

In the June 2026 edition, Baby products was also lowest, at 3.82. On the trend base the 18-category mean was 4.01 in both editions (§1.1).



Mean AI trust across 18 categories for all 1,493 shoppers and for the 1,254 who differentiate between categories, with the 239 who refuse AI in every category removed.

(AI trust by category: n=1,493 full sample; refusers n=239; differentiators n=1,254; source: §10.1.)

10.2 Buyers of baby products use AI and delegate to it the most

149 respondents bought baby products or child-safety items in the past 12 months. They use AI to research purchases at 58.4%, against 36.8% of the full sample. Among the 87 who use AI, 36.8% have completed an AI purchase, against 21.6% of AI users overall. They rate AI trust in the baby category at 4.70 out of 7, against 3.56 among everyone who did not buy in it.

The gap is not an age effect. Within every age band from 18–24 through 45–54, buyers of baby products complete AI purchases at roughly twice the rate of other shoppers their age. The two oldest bands hold fewer than ten baby-product buyers each and are not reported.

These figures rest on 149 buyers, 87 AI users, and 32 who have completed an AI purchase. They are directional, and Edition 3 will include a deliberate oversample of parents so they can be stated with precision.

(Baby-product buyers: n=149; AI users n=87; completers n=32; directional; source: §10.2.)

10.3 Younger, male, and higher-income shoppers use AI more

Adoption by age is in §2.1. Verification and completion by age among AI users:

AGE	AI USERS	VERIFY	COMPLETED AN AI PURCHASE
18–24	83	94.0%	25.3%
25–34	116	94.0%	33.6%
35–44	118	86.4%	21.2%
45–54	77	93.5%	22.1%
55–64	83	83.1%	14.5%
65 or older	73	86.3%	6.8%

Men use AI for product research more than women (46.0% against 37.2% on the trend base; 50.1% against 35.0% in the June 2026 edition). Among AI users, men and women verify at similar rates (90.8% and 88.2%) and complete AI purchases at similar rates (22.4% and 20.8%). The difference is over payment. 50.2% of women would refuse an agent any payment access against 36.9% of men, and 42.0% of women would let an agent spend nothing against 33.6% of men. 49.1% of women and 43.0% of men cap an unchecked AI recommendation under \$25.

Adoption rises with income, from 30.3% among households under \$30,000 to 47.2% at \$100,000–\$150,000.

(Age cells 73 to 118 AI users; gender: men n=295 AI users of 744, women n=255 of 749; income bands 70 to 512; source: §10.3.)

10.4 Store-only shoppers use AI less and trust retailers less

The June 2026 edition included only shoppers who buy online. Edition 2 also kept shoppers who buy only in physical stores, as a separate segment of 246 respondents, 16.5% of complete responses, and they never enter a trend number. They use AI to research purchases at 13.4%, rate confidence in their last purchase at 6.12 out of 10, and score retailers a Net Retailer Trust Score of –52.8, against 41.5%, 7.03, and –23.7 among online shoppers.

(Store-only segment: n=246; source: §10.4.)

10.5 The West leads adoption and the Northeast checks the most

Four Census regions were assigned from passive geolocation. 1,008 of 1,493 complete responses could be located, so regional figures are directional.

REGION	N	USED AI FOR RESEARCH	VERIFY (OF AI USERS)	CAP UNDER \$25	COMPLETED AN AI PURCHASE (OF AI USERS)
West	174	40.8%	85.9%	47.7%	25.4%
South	422	37.2%	91.1%	45.3%	23.6%
Northeast	213	33.3%	93.0%	53.5%	12.7%
Midwest	199	32.2%	90.6%	48.2%	26.6%

The West leads adoption. The Northeast trails on adoption and completion and verifies the most. AI-user cells run from about 64 to 157.

(Regions: located subsample n=1,008; directional; source: §10.5.)

Reading the numbers

This section is interpretation. Everything above it reports what the survey measured. What follows is how Product.ai Research reads those measurements, with each reading tied to the figure it rests on.

11.1 Reading one: The shoppers who let AI buy are the ones who check it most

The two headline facts of this edition look like a contradiction. Adoption is flat, AI ranks last among research sources, and nine in ten AI users check its answers. Yet one in five AI users has let an assistant buy something, and two in three have started to. Completion rises with verification habit, from 8.8% among shoppers who never check to 28.8% among those who always do (§5.3).

The natural assumption is that the people handing an assistant their card would be the least skeptical of AI answers. The data point the other way. The 119 shoppers who have let an assistant complete a purchase are also the most permissive in the survey. 42.9% would let an agent order without asking and 54.6% would give it unlimited payment access, against 16.9% and 19.8% of AI users at large. These findings suggest that checking AI and trusting it to buy go together.

11.2 Reading two: One in six shoppers refuses AI, and the rest trust it moderately

The averages in §10.1 combine two groups. One shopper in six refuses AI in all 18 categories and, in 95% of cases, had not used AI to research a purchase in the past 90 days. Among the other five in six, AI's average trust is above the scale's neutral midpoint (4 of 7) in every one of the 18 categories, from 4.18 to 4.55 (§10.1). "Consumers distrust AI shopping" is true of the first group and false of the second, and the average describes neither.

The split, more than the average, is the durable finding of this edition. The question for Edition 3 is whether the share of shoppers who distrust AI in every category shrinks or holds.

11.3 Reading three: Shoppers trust the assistants they already use

Gemini and ChatGPT lead both usage and the most-trusted ranking, within two points of each other on trust. Claude and Perplexity draw the largest least-trusted shares and are used by 17.6% and 11.3% of AI users. The order of trust matches the order of usage for six of seven assistants (§3.2), and roughly half of shoppers declined to pick most assistants either way.

The ranking tracks familiarity as much as trust. An assistant a shopper has not used tends to end up in the least-trusted column. That is the cold-start problem every new shopping agent faces, where nobody trusts what nobody has used yet, and it is why this report publishes the ranking and its critique together.

11.4 Reading four: Older shoppers punish hidden sponsorship far more than younger ones

Across all AI users, a hidden sponsorship costs 10.5 points more trust than a disclosed one (§8.2). Shoppers 55 and older penalize the hidden one by more than 20 points, while shoppers aged 25–34 mind the label slightly more than the concealment. These age groups hold 73 to 118 AI users each, so the split is directional. For assistants that show sponsored results, these findings imply that older shoppers weigh disclosure most heavily.

11.5 Reading five: Shoppers believe AI on specifications more than the evidence supports

61.1% of AI users believe AI answers are usually accurate about product specifications (§8.3). Product.ai’s [AI Shopping Divergence Study](#), which asked four assistants the same shopping questions through their developer APIs, found a confirmed conflict on 75% of plain spec and fact lookups, the question type on which the assistants agreed most. These findings suggest that checking is warranted even on the answers shoppers trust most.

What Edition 3 will add and test

12.1 Four additions planned for Edition 3

1. **A parent oversample.** Buyers of baby products are the most likely to let an AI assistant complete a purchase (36.8% of those who use AI, against 21.6% overall), on a base of 149 buyers and 87 AI users (§10.2). Edition 3 oversamples households with young children so the finding can be stated with precision or retired.
2. **Grok and Meta AI on the assistant list.** Both recurred as write-ins in Edition 2 (§3.1).
3. **Closed options for the trust-reasons question.** 54% of answers to the open question in §3.3 could not be coded. Edition 3 offers structured reasons and reserves open text for the “verified” definition, which produced usable answers.
4. **A definition of completion inside the question.** Edition 3 asks separately whether the assistant placed the order and charged payment, so that “completed a purchase” no longer depends on the respondent’s reading (§5).

12.2 Five open questions for Edition 3

1. **Will the share of shoppers who double-check AI recommendations keep rising?** 85.7% in the first edition, 89.9% in this one. Edition 3 re-measures on the same base.
2. **Will more AI users let an assistant complete a purchase, and will those who always verify still lead?** 21.6% of AI users in this edition; 28.8% of those who always verify.
3. **Will the share of shoppers who distrust AI in every category shrink?** 16.0% of shoppers refuse AI in every category (§10.1).
4. **Does the Disclosure Premium change as sponsored results become more common?** 10.5 points in this edition, with 41.5% of AI users having noticed a sponsored recommendation.
5. **Will the store-only segment start to use AI?** 13.4% adoption in this edition.

12.3 How often the survey runs and how editions compare

Product.ai holds question wording, response options, and category lists constant for every tracked measure, and fields each edition on the same general-population panel with the same qualifying conditions. Edition 3 is planned for spring 2027. Each repeated question appears on this page as a time-series chart that gains a point per edition. Edition 1's full text stays at its own address, and the data for every edition stays public.

12.4 New baselines set in Edition 2

- **Completed AI purchases:** 21.6% of AI users; 66.9% funnel touch.
- **Agent Allowance:** 37.5% allow nothing; 61.0% allow under \$25 or nothing.
- **Payment access:** 43.1% refuse; 19.8% grant without limits.
- **Disclosure Premium:** 10.5 points.
- **Believed accuracy of specifications:** 61.1%.
- **Awareness that AI can repeat marketing as fact:** 75.1% knew or suspected.
- **The “verified” gap:** 14.0% describe independent checking; 14.9% do not know.
- **Multi-assistant use:** 64.0%; 2.25 assistants on average.
- **Store-only segment:** 13.4% adoption; Net Retailer Trust Score –52.8.
- **Holiday behavior:** 36.5% of AI users had used an assistant for holiday shopping by late August.

Limitations

13.1 Two editions do not make a trend

Two editions make a comparison, and a trend needs more than two points. Each held-steady reading in §1 says that a measure did not move by more than the pre-registered threshold between the two editions; it does not say the measure will not move. Causal claims are not supported by the design.

13.2 U.S. residents only

Both editions sample U.S. residents aged 18 and over. Findings do not generalize to other markets.

13.3 Most findings describe the 550 AI users, not all 1,493 shoppers

Verification, completion, permissions, platform, and disclosure findings describe the 550 respondents who used AI to research a purchase in the past 90 days (517 on the trend base), not all shoppers. Every such figure states its base at the point of citation.

13.4 All behavior is self-reported

All behavior is self-reported. Verification is a socially approved answer and may be over-reported; the direction of the finding is more reliable than its precise level. “Completed a purchase” is the respondent’s reading of the question and may include a checkout the assistant opened and the shopper confirmed inside it (§5).

13.5 Some subgroups are too small for precise figures

Regional AI-user cells (64 to 157), age bands among AI users (73 to 118), buyers of baby products (149; 87 AI users; 32 completers), the 119 shoppers who have completed an AI purchase, and the assistant shares among them are directional. The report marks each as such where it appears.

13.6 The platform trust ranking cannot yet be recomputed from the public data

The most- and least-trusted grid in §3.2 is drawn from the survey platform’s standard report because the combined grid is not included in the raw-data export. The ranking can be cited; it cannot yet be recomputed or cut by segment from the public dataset. Product.ai will add the raw columns to the dataset when they are supplied.

13.7 Region is known for two-thirds of respondents

Region was assigned by passive geolocation for 1,008 of 1,493 complete responses. Unlocated respondents are missing at random only by assumption.

13.8 Neither edition is weighted to Census distributions

Neither edition is weighted to Census distributions. Panel quotas balanced age, gender, and region. The pre-registered change threshold of about 4 points was set to absorb the design effect of that choice. Census post-stratification remains planned.

Methodology

Sample frame

Edition 2 sampled U.S. residents aged 18 and over who shop for products, online or in stores, through a general-population panel. Edition 1 required online shopping to qualify. Edition 2 kept store-only shoppers as a tagged segment (n=246) and computes every edition-over-edition comparison on the online-shopper base (n=1,247) so the two editions compare like with like.

Fielding details

PARAMETER	EDITION 2	EDITION 1
Field window	August 27 to September 1, 2026	April 27, 2026
Survey platform	Alchemer	Alchemer
Sample panel	General-population panel	General-population panel
Instrument	43 questions; 12 questions from the June 2026 edition repeated word for word	35 questions
Total records	2,624	3,638
Disqualified at screen	760	1,778
Partial completions	371	397
Complete responses	1,493	1,463
Trend base	Online shoppers, n=1,247	n=1,463
AI-user base	550 (517 on the trend base)	623
Median completion time	7.2 minutes	not reported
Analysis basis	Complete responses only	Complete responses only

Change threshold

Before fielding, the research team set the rule for reporting an edition-over-edition change. A move of about 4 points or more, after design effects, is reported as a change, and smaller moves are reported as held steady. The rule was written into the analysis plan and applied to every measure in §1.

Screener and routing

Respondents under 18 or outside the United States were disqualified. Store-only shoppers continued through the survey as a segment. Respondents who had not used an AI assistant to research a purchase in the past 90 days skipped the AI-user blocks (verification, permissions, platforms, disclosure, completion, holiday) and continued to demographics. Three questions asked of every respondent in Edition 1 (the AI-trust grid, the \$25 threshold, and review trust) were placed ahead of that gate so they remained comparable and could be cut by AI use.

Sample composition

AGE	N	SHARE
18-24	153	10.2%
25-34	234	15.7%
35-44	269	18.0%
45-54	238	15.9%
55-64	271	18.2%
65 or older	328	22.0%

GENDER	N	SHARE
Female	749	50.2%
Male	744	49.8%

HOUSEHOLD INCOME	N	SHARE
Under \$30,000	512	34.3%
\$30,000-\$59,999	429	28.7%
\$60,000-\$100,000	323	21.6%
Over \$100,000	229	15.3%

ETHNICITY	N	SHARE
White, non-Hispanic	870	58.3%
Hispanic or Latino	272	18.2%
Black or African American	186	12.5%
Asian or Pacific Islander	106	7.1%
Other	59	4.0%

98.9% of complete responses geolocated to the United States.

Denominators

Every figure in this report is paired with the base it was computed on. The four bases used most:

- **Full sample, n=1,493.** All complete responses. Used for the “verified” open text, the AI-trust grid, and full-sample descriptive figures.
- **Trend base, n=1,247.** Complete responses from shoppers who buy online at least some of the time. Used for every comparison with Edition 1 and for the trust hierarchy, PRCl, retailer trust, and review trust headline figures.
- **AI users, n=550.** Complete responses from shoppers who used an AI assistant to research a purchase in the past 90 days. Used for verification, permissions, platforms, disclosure, completion, and holiday

findings. On the trend base the AI-user count is 517.

- **Routed sub-bases.** Verifiers (n=493), error-finders (n=443), shoppers who limit or refuse payment access (n=441), completers (n=119), buyers of baby products (n=149), store-only shoppers (n=246), and the located subsample (n=1,008). Each is stated at the point of citation.

Scale encoding

The survey platform exports rating scales as a mix of text labels and numbers. Every mean in this report was computed after mapping each label to its numeric anchor (for example, “Extremely confident” to the scale maximum and “Neutral” to the midpoint). In both editions, every 1–7 grid is averaged over every answer, with the worded anchors mapped to 1, 4, and 7. Scales: the last-purchase confidence and retailer trust questions run 0–10; the trust hierarchy, category confidence, AI-trust-by-category, review accuracy, believed accuracy, bias, and allowance-comfort questions run 1–7.

Data sources

Every figure in this report except one block was recomputed from the raw response export. A vetted copy of the responses, with identifying fields removed, is published alongside the report. The most- and least-trusted grid (§3.2) is taken from the survey platform’s standard report because the combined grid is not included as raw columns; the raw columns will be added to the public dataset when supplied.

Quality controls

Respondents failing the age or residency screen were disqualified at intake. Only complete responses entered the analysis. Routing to the AI-user blocks was enforced at the instrument level. Completion times were reviewed; the median was 7.2 minutes.

Weighting

Neither edition is weighted to Census distributions. Panel quotas balanced age, gender, and region. Census post-stratification is planned, and the first re-weighted comparison will re-weight both editions to the same profile.

Footnote and citation convention

Every figure has an in-line footnote in the form (*value; base; source: §anchor*). Worked example for the headline verification figure:

89.9% of U.S. online shoppers who used an AI assistant for product research verified its recommendation through another source before buying. (*n=517 AI users among 1,247 online shoppers on the trend base; 41.0% always verify + 48.9% sometimes verify; Edition 1: 85.7% of 623; source: §2.2*)

Conflict of interest

Product.ai builds verification products for commerce. The “verified” question (§9.1) follows the same reporting rules as every other figure.

Artifacts

Edition 2's dataset (PII-stripped microdata, CSV), the codebook, and the report PDF are published under CC BY 4.0 at edition-named addresses. Edition 1's dataset and PDF remain at their original addresses.