

The 2026 Trust in AI Commerce Report

Among the 43% of U.S. online shoppers who used AI for product research in the past 90 days, 86% verified the AI's recommendation through another source before buying. The verification layer is the missing infrastructure.

BY MICHAEL QUOC, DAKOTA NUNLEY, ELENA MADRIGAL, SEAN FISHER · PRODUCT.AI RESEARCH · JUNE 23, 2026

Key Takeaways

- 1. **86% of U.S. online shoppers who used AI for product research verified the AI's recommendation through another source before buying.** 45% always verified; 41% sometimes verified; 14% trusted the recommendation without verifying. (*n=623 AI users among 1,463 surveyed; source: §1.3*)
 - 2. **42% of U.S. online shoppers would not trust an AI recommendation for any purchase over \$25 without checking another source first.** 61% cap at \$50 or under. 5% would trust AI for a purchase over \$500. (*n=1,463; unrounded value 41.8%; source: §2.1*)
 - 3. **The Purchase Research Confidence Index (PRCI) on shoppers' most recent \$50+ online purchase measures 6.95 out of 10.** Across 18 product categories, per-category research confidence averages approximately 4.98 and ranges from 4.49 (Baby Products) to 5.08 (Apparel). (*n=1,453; source: §3.2, §3.3; see Methodology – Mixed-format encoding correction for context on the corrected value*)
 - 4. **Among seven information sources used in product research, AI tools rank sixth at 4.51 out of 10.** Friends and family rank first at 5.13. No source measures above 5.2; the full range spans 0.65 points. (*n=1,463; source: §4.1*)
 - 5. **AI adoption for product research drops between the 35–44 and 45–54 age cohorts.** 18–24: 56% / 25–34: 58% / 35–44: 58% / 45–54: 44% / 55–64: 30% / 65+: 21%. Under-44 cohorts at 55–58%; 55+ cohorts under 30%. (*n=1,463; source: §5.2*)
 - 6. **Apparel sits at opposite ends of two indices: highest for per-category research confidence (5.08) and lowest for per-category AI trust (4.30) across 18 categories surveyed.** (*n=1,463; source: §3.3, §5.1*)
-

Executive Summary

The Trust in AI Commerce Report measures U.S. online-shopper behavior, trust, and verification practice around AI product recommendations. Conducted in April 2026 against a general-population panel of 1,463 U.S. online shoppers, the survey covered AI shopping behavior, purchase research confidence across 18 product categories, the broader information trust hierarchy, retailer and review trust, and category-level patterns by age, gender, income, and ethnicity. The headline finding: among the 43% of respondents who used AI for product research in the past 90 days, 86% verified the AI's recommendation through another source before buying. 14% trusted the recommendation without verifying. Across the full surveyed group, 42% would not trust an AI recommendation for any purchase over \$25 without checking another source first.

The verification pattern sits inside a broader trust landscape the research also measured. The Purchase Research Confidence Index (PRCI) — a confidence framework with two sub-measures, both reported here — reads at 6.95 out of 10 for shoppers' most recent online purchase of \$50 or more. Across 18 product categories, per-category research confidence ranges from 4.49 (Baby Products) to 5.08 (Apparel), with a cross-category mean of approximately 4.98. The information trust hierarchy ranks seven product-research sources from most to least trusted. Friends and family lead at 5.13 out of 10; AI tools rank sixth at 4.51. No source measures above 5.2 on the scale; the full hierarchy spans 0.65 points across all seven sources. Net retailer trust on the statement "Online retailers always have my best interest as a customer in mind" measures -25 NPS-style (mean 6.35 out of 10).

The findings describe a U.S. online-shopper population in which AI for product research has reached 43% past-90-day usage and 20% per-purchase usage. Among AI users, 86% verify the AI's recommendation through another source before buying. The information sources shoppers use for product research overall sit on a hierarchy where no source measures above 5.2 out of 10. The research did not separately measure which sources shoppers consult specifically to verify AI recommendations; that question, along with re-measurement of the canonical findings above, is documented in §6 as part of the open research questions and planned additions to subsequent rounds.

AI Shopping Behavior & The Verification Pattern

The research measured four AI shopping behaviors: whether shoppers used AI to research a purchase, what other tools they used for product research, whether they cross-checked the AI's recommendation against another source before buying, and how they felt about the purchase

afterward. The 86% verification finding in §1.3 uses the AI-user subgroup as its denominator; the other measures describe the full surveyed group of 1,463.

1.1 How often U.S. online shoppers use AI for product research

43% of U.S. online shoppers used AI for product research in the past 90 days (unrounded 42.6%; based on responses from 1,463 U.S. online shoppers). The 43% includes any use of an AI assistant — ChatGPT, Gemini, Claude, Perplexity, or a comparable tool — to research a product during that window. A respondent who used AI once for a single product and a respondent who used AI on every \$50+ purchase both count toward the 43%.

The 43% reflects the average across all age groups. The age-level rate varies: 56% of 18–24-year-olds, 58% of 25–34s, 58% of 35–44s, 44% of 45–54s, 30% of 55–64s, and 21% of those 65 and older. The age-level breakdown is reported in §5.2.

(43% AI usage past 90 days: n=1,463; unrounded 42.6%; source: Methodology — Screener logic. Age decomposition: §5.2.)

1.2 How often shoppers used AI to research their most recent purchase

20% of U.S. online shoppers used AI to research their most recent online purchase of \$50 or more (unrounded 20.3%; n=1,463). This measure is narrower than the past-90-day usage figure in §1.1. The past-90-day question asks whether the respondent used AI for product research at any point during that window, on any product. The 20% figure here asks whether the respondent used AI on one specific purchase — the most recent one over \$50. A respondent who used AI three months ago for a different product but did not use AI on the most recent purchase over \$50 would count toward the 43% in §1.1 but not toward the 20% here.

The category mix of "most recent \$50+ purchases" is reported in §3.2, where Apparel (30% of last-purchase mentions) and Electronics (20%) lead. Per-category AI trust — how much shoppers trust AI's recommendations within each category — is reported in §5.1.

(20% AI usage on most recent \$50+ purchase: n=1,463; unrounded 20.3%; source: Methodology — Fielding details. Category mix: §3.2.)

1.3 What AI users do after they receive an AI recommendation

Among the 623 respondents who used AI for product research in the past 90 days, **86% verified the AI's recommendation through another source before completing the purchase** (unrounded 85.6%). Verification was self-reported and described in the question as cross-checking the AI's answer against an additional source — for example, looking the recommended

product up on the brand's website, reading customer reviews, asking a friend, watching a video review, or running a Google search before buying.

The 86% breaks down as follows: **45% of AI users said they always verify AI recommendations** through another source (277 of 623). **41% said they sometimes verify** (257 of 623). **14% said they trust AI recommendations without verifying** through another source (89 of 623).

The 86% figure describes only the 623 people who used AI for product research. It does not describe the full 1,463 surveyed. A respondent who never used AI for product research could not have verified an AI recommendation, because they never received one. Applying the full-sample figure would mix AI users with people who never asked AI for a recommendation in the first place. The full denominator rationale is documented in Methodology – Denominator definitions (AI-user subsample).

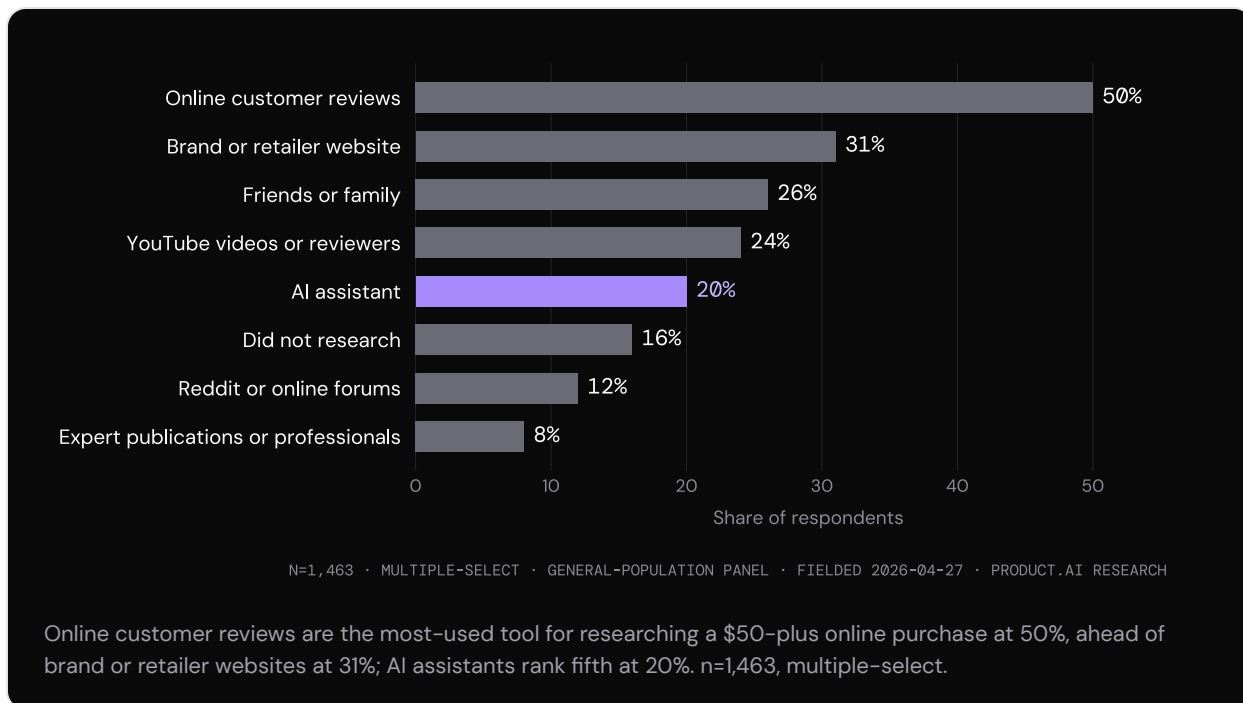
Among AI users, satisfaction with the AI tool itself measured **5.46 out of 7** on average (623 AI users; 1–7 scale). **32.9% rated their AI satisfaction at the top of the 7-point scale** ("Extremely satisfied").

(86% verification: n=623 AI users among 1,463 surveyed; 45% always verified + 41% sometimes verified; source: §1.3; denominator anchor: Methodology – Denominator definitions (AI-user subsample). AI satisfaction mean 5.46/7: n=623 AI users.)

1.4 What tools shoppers used to research their most recent purchase

Before completing their most recent \$50+ online purchase, respondents were asked which tools they consulted for product research. Respondents could pick more than one. The distribution across all 1,463 respondents:

RESEARCH TOOL	SHARE OF RESPONDENTS
Online customer reviews (Amazon, retailer sites)	50%
Brand or retailer website	31%
Friends or family	26%
YouTube videos or reviewers	24%
AI assistant (ChatGPT, Gemini, Claude, Perplexity)	20%
Did not research	16%
Reddit or online forums	12%
Expert publications or professionals	8%



Online customer reviews on Amazon and retailer sites place first at 50%. Brand or retailer websites place second at 31%. Friends and family place third at 26%. AI assistants rank fifth at 20%. Reddit forums rank seventh at 12%, and expert publications eighth at 8%. 16% reported doing no research before their most recent \$50+ purchase.

The research did not include a separate question on which sources shoppers consult when they verify an AI recommendation specifically. The table above describes the tools shoppers use for product research broadly. A dedicated verification-source question is documented in §6 as a question planned for future research.

(Research tool stack: n=1,463; respondents could pick more than one tool; source: §1.4.)

1.5 What shoppers reported after the purchase

After completing their most recent \$50+ online purchase, respondents reported how they felt about the decision. Respondents could register more than one feeling about the same purchase. The distribution across all 1,463 respondents:

POST-PURCHASE FEELING	SHARE	RESPONDENTS
Felt confident I made the right choice	77%	1,131
Wished I had researched more	14.6%	214
Found a better option afterward	12.5%	183
Returned the product	4%	57
None of the above	5%	69

77% of respondents felt confident in the decision they made on their most recent \$50+ online purchase. Confidence is the most-reported outcome.

Two other responses also registered: **14.6% wished they had researched more before buying,** and **12.5% found a better option afterward.** Because respondents could register both feelings on the same purchase, those two responses overlap. The sum of the two response counts (14.6% + 12.5%) approximates 28% of all respondents reporting at least one of the two forms of post-purchase reflection, with overlap between them allowed.

Both rates are quantifiable and will be re-measured by future research on the same questions. Whether either rate moves between rounds is one of the open research questions documented in §6.

(Post-purchase outcomes: n=1,463; respondents could register more than one feeling; source: §1.5.)

1.6 Summary of measured findings

The four behaviors measured and reported in §1.1 through §1.5:

- **43%** of U.S. online shoppers used AI for product research in the past 90 days (§1.1).
- **20%** used AI to research their most recent online purchase of \$50 or more (§1.2).
- Among the 623 respondents who used AI for product research, **86%** verified the AI's recommendation through another source before buying; **14%** did not (§1.3).
- Across all 1,463 respondents, the most-used product-research tools are online customer reviews (50%), brand or retailer websites (31%), and friends or family (26%); AI assistants rank fifth at 20% (§1.4).
- **77%** of respondents felt confident in their most recent \$50+ purchase decision; **14.6%** reported wishing they had researched more; **12.5%** reported finding a better option afterward (§1.5).

The research did not include a separate question on which tools shoppers consult specifically when they verify an AI product recommendation. That measurement is documented in §6 as a

question planned for future research.

The \$25 AI Autonomy Threshold

Respondents were asked the maximum dollar amount they would spend on a product recommended by an AI assistant without consulting any other source. The distribution of answers is the AI Autonomy Threshold.

2.1 The \$25 threshold finding

42% of U.S. online shoppers said they would not trust an AI recommendation for any purchase over \$25 without checking another source first (unrounded 41.8%; 611 of 1,463 respondents). The question asked respondents to state the maximum dollar amount they would spend on a product recommended by an AI assistant without consulting any other source. 42% picked the lowest available bracket – under \$25.

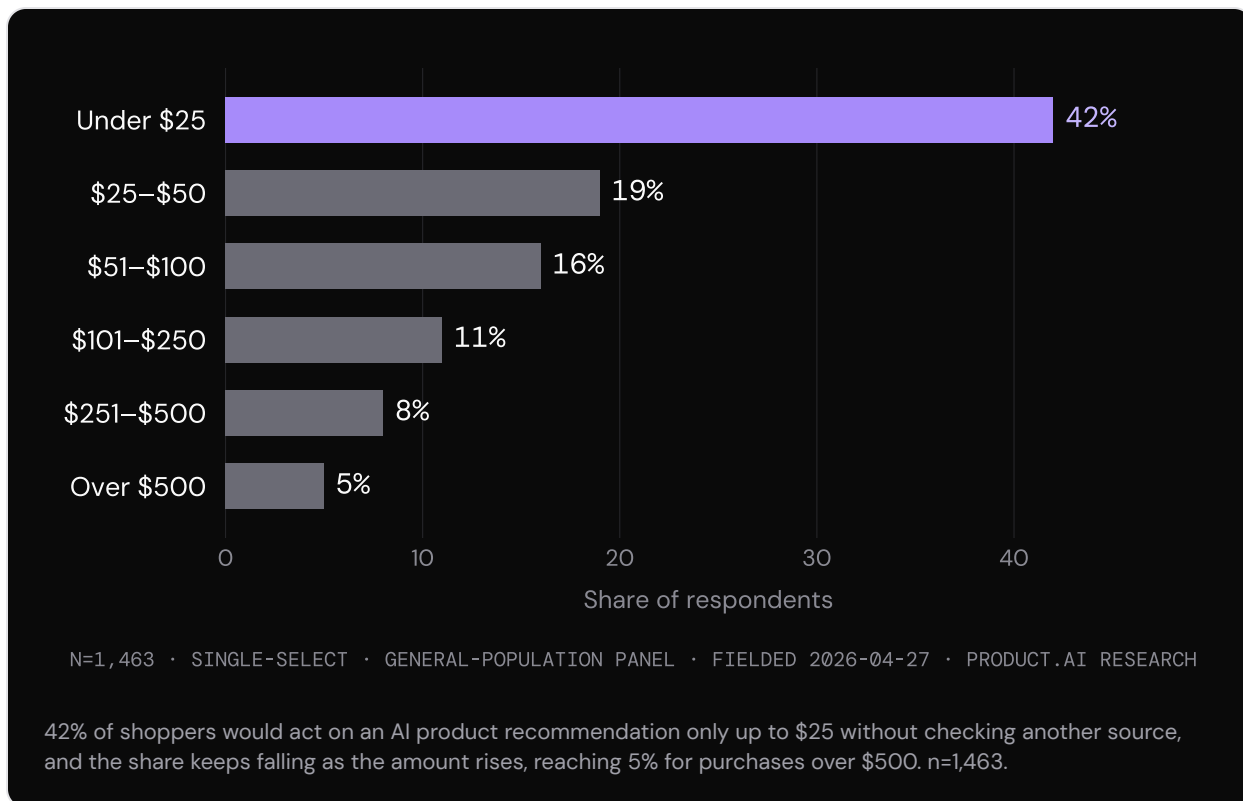
The \$25 figure is the largest single response on the question. The next-most-common response is the adjacent \$25–\$50 bracket at 19%.

(42% under \$25; n=1,463; unrounded 41.8%; source: §2.1; methodology integrity-flag: §1.6a.)

2.2 The full dollar-tier autonomy curve

Respondents were given six dollar brackets to choose from. The distribution across all 1,463 respondents:

MAXIMUM SPEND ON AI RECOMMENDATION WITHOUT VERIFYING	SHARE	RESPONDENTS
Under \$25	42%	611
\$25–\$50	19%	278
\$51–\$100	16%	233
\$101–\$250	11%	153
\$251–\$500	8%	115
Over \$500	5%	72



60.8% of respondents cap their unverified AI trust at \$50 or less. 4.9% would trust AI for any purchase over \$500. The remaining 34% sit between \$50 and \$500.

The distribution descends bracket by bracket. Each higher price tier holds a smaller share of respondents than the one below it.

(Dollar-tier distribution: n=1,463; one respondent did not answer; source: §2.2.)

2.3 The curve restated cumulatively

Restated cumulatively — the share of respondents whose threshold is at or above a given price point — the same data reads:

PRICE POINT	SHARE WILLING TO TRUST AI WITHOUT VERIFYING AT OR ABOVE THIS AMOUNT
\$25	58% (852 of 1,463)
\$50	39% (574)
\$100	23% (340)
\$250	12% (187)
\$500	5% (72)

For comparison, §1.2 reports that 20% of U.S. online shoppers used AI to research their most recent online purchase of \$50 or more. The dollar-tier curve here describes a related but distinct measure — what shoppers say they would trust AI for hypothetically, not what they actually did on a specific purchase.

(Cumulative restatement of the §2.2 distribution; n=1,463.)

2.4 What this measure tracks across rounds

The dollar-tier question uses six fixed brackets. Future research using the same question produces a new distribution that can be compared to this one bracket by bracket. Bracket boundaries are the same regardless of when the question is asked. Whether the share of respondents capping AI trust under \$25 moves up or down the price scale, and where the highest-trust 5% cohort settles, are both quantifiable comparisons.

The thresholds and their movement across rounds are listed in §6 as open research questions.

Purchase Research Confidence Index (PRCI)

The Purchase Research Confidence Index (PRCI) is a confidence framework with two sub-measures, both reported here. PC-1 captures confidence in the respondent's most recent significant online purchase. PC-6 captures general confidence shopping across 18 product categories. The two are distinct sub-questions within the same framework on the same survey.

3.1 What the PRCI measures

The PRCI contains two sub-measures, both reported in the research:

(a) Last-purchase confidence (PC-1). A single retrospective rating on a 0–10 scale, asking the respondent how confident they were that they found the best product before completing their most recent online purchase of \$50 or more.

(b) Per-category general confidence (PC-6). A set of 18 ratings on a 0–10 scale, asking the respondent how confident they generally feel about finding the best product when shopping for each of 18 product categories.

PC-1 asks about a specific purchase the respondent already completed. PC-6 asks about general shopping confidence across categories, regardless of whether the respondent has bought in those categories recently. Both anchor on the same 0–10 scale.

3.2 Last-purchase confidence: 6.95 out of 10

The mean confidence rating for shoppers' most recent online purchase of \$50 or more was **6.95 out of 10** (n=1,453). **25%** of respondents felt **fully confident** in their decision (rated 10, "Extremely confident"). **46.7%** rated their confidence at **7 or below**.

The categories that made up the "most recent \$50+ purchase" mix:

CATEGORY	SHARE OF MOST-RECENT-PURCHASE RESPONSES
Apparel (shoes, clothing, accessories)	30%
Electronics (smart phone, computer, tablet)	20%
Grocery / food delivery	8%
Video games	7%
All other 14 categories	balance

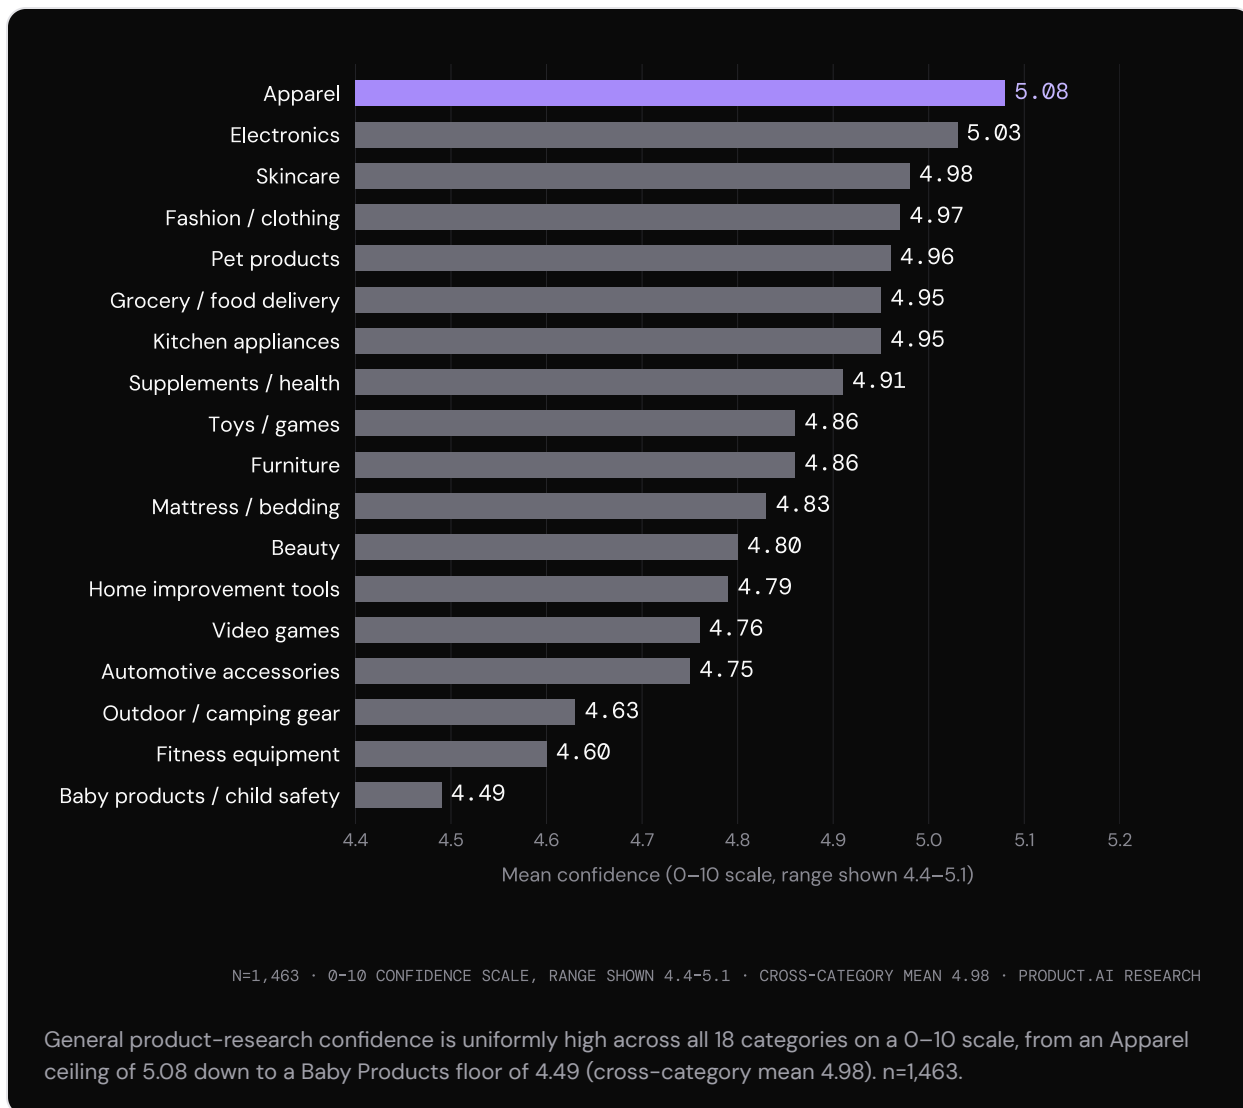
The 6.95 mean reflects the population mix above.

(Last-purchase confidence mean 6.95/10: n=1,453; source: §3.2; see Methodology – Mixed-format encoding correction.)

3.3 Per-category general confidence: 18 categories

The 18-category general confidence question (PC-6) was answered by all 1,463 respondents on a 0–10 scale. The cross-category mean across all 18 categories is approximately **4.98**. Per-category means range from 4.49 (lowest) to 5.08 (highest).

CATEGORY	MEAN (0-10)	VS. CROSS-CATEGORY AVG
Apparel (shoes, clothing, accessories)	5.08	+0.10 (ceiling)
Electronics (smart phone, computer, tablet)	5.03	+0.05
Skincare (lotion, supplements, sunscreen)	4.98	0.00
Fashion / clothing (jacket, jeans, handbag)	4.97	-0.01
Pet products	4.96	-0.02
Grocery / food delivery	4.95	-0.03
Kitchen appliances	4.95	-0.03
Supplements / health products	4.91	-0.07
Toys / games	4.86	-0.12
Furniture	4.86	-0.12
Mattress / bedding	4.83	-0.16
Beauty (makeup, cosmetics)	4.80	-0.18
Home improvement tools	4.79	-0.19
Video games	4.76	-0.22
Automotive accessories	4.75	-0.23
Outdoor / camping gear	4.63	-0.35
Fitness equipment	4.60	-0.38
Baby products / child safety	4.49	-0.49 (floor)



The full 18-category range spans 0.59 points (4.49 to 5.08). The category-level ceiling sits below 5.5 on the 0–10 scale.

(Per-category confidence 18-category table: n=1,463; source: §3.3.)

3.4 The floor — Baby products at 4.49

Baby products / child safety reports the lowest mean per-category confidence rating at **4.49 out of 10**, 0.49 points below the cross-category mean. The category covers strollers, car seats, baby gates, and similar safety-relevant items.

How baby-products confidence varies by demographic is reported in §5.

3.5 The ceiling — Apparel at 5.08

Apparel (shoes, clothing, accessories) reports the highest mean per-category confidence rating at **5.08 out of 10**, 0.10 points above the cross-category mean. Apparel is also the largest share of "most recent \$50+ purchase" responses (30%, per §3.2).

Apparel sits at the top of the PRCI per-category index (5.08) and at the bottom of the Category-Specific AI Trust Index (4.30, per §5.1) — the same category appears at opposite ends of the two indices. The cross-index observation is reported in §5.1.

3.6 What this index tracks across rounds

Both PRCI sub-measures (PC-1 and PC-6) use fixed scales and identical question wording. Future research using the same questions produces values directly comparable to this round. Changes in the last-purchase confidence mean, the per-category cross-category mean, the ceiling category, the floor category, or any specific per-category mean are quantifiable across rounds.

The PRCI is one of the named indices the research will track. Open questions about the index are listed in §6.

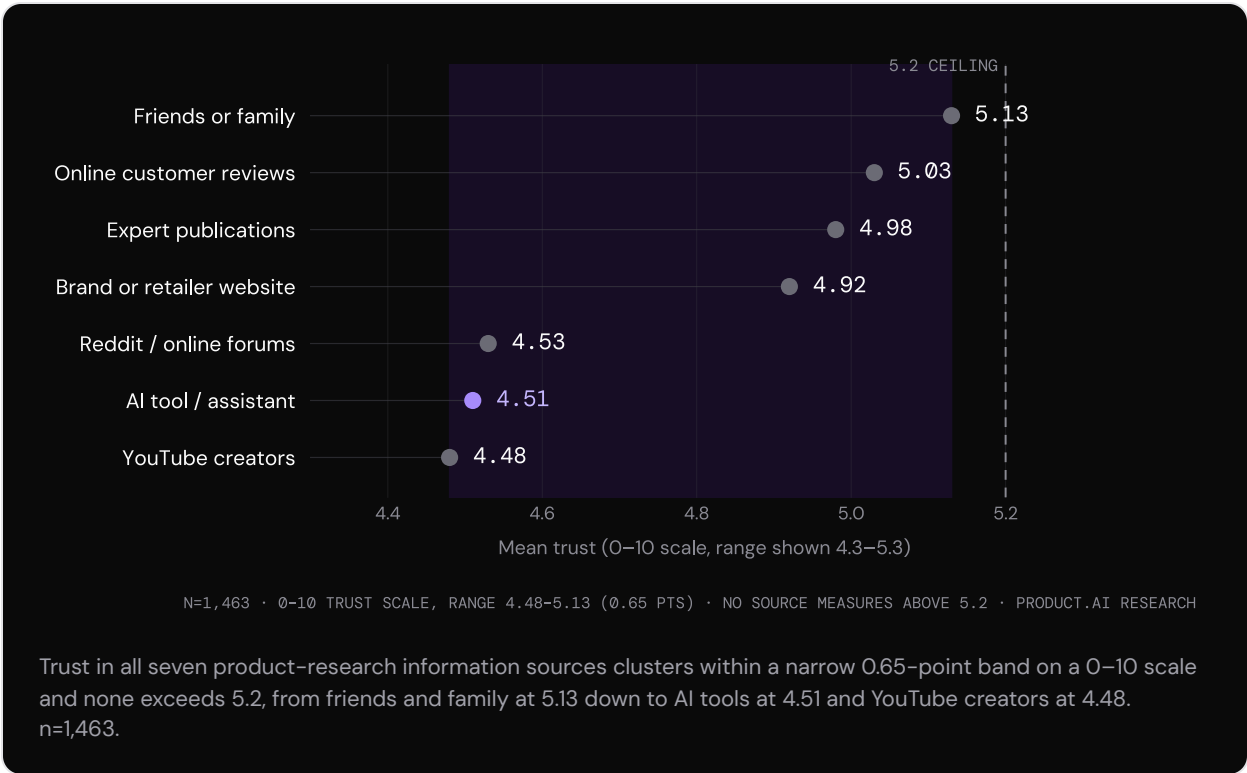
The Information Trust Hierarchy

Two related sets of measures sit in this chapter. The information trust hierarchy measures how much U.S. online shoppers trust each of seven product-research sources on a 0–10 scale. The broader trust landscape around AI shopping adds retailer trust and review trust.

4.1 Seven information sources ranked by mean trust

Respondents rated their trust in each of seven sources when researching a product they might buy online. The 1,463-respondent mean for each source on the 0–10 scale:

INFORMATION SOURCE	MEAN TRUST (0-10)	RANK
Friends or family	5.13	#1
Online customer reviews	5.03	#2
Expert publications or professionals	4.98	#3
Brand or retailer website	4.92	#4
Reddit or online forums	4.53	#5
AI tool / assistant	4.51	#6
YouTube creators or reviewers	4.48	#7



Friends and family rank first at 5.13. AI tools rank sixth at 4.51. YouTube ranks seventh at 4.48.

(Information trust hierarchy: n=1,463; source: §4.1.)

4.2 The 0.65-point band

The full range across all seven sources spans 0.65 points (from 4.48 at the bottom to 5.13 at the top). **No source measures above 5.2** on the 0-10 scale.

The four highest-ranked sources (Friends/Family, Customer Reviews, Expert Publications, Brand/Retailer Website) sit within a 0.21-point band of each other (4.92-5.13). The three lowest-

ranked sources (Reddit, AI, YouTube) sit within a 0.05–point band of each other (4.48–4.53). AI's position (4.51) sits 0.62 points below the top source and 0.03 points above the bottom source.

(Range and band measurements derived from the §4.1 means; n=1,463.)

4.3 Reading the range

Two ways to read the trust hierarchy describe the same data:

- .. **By rank order:** AI ranks sixth of seven. Friends and family lead. Reddit, AI, and YouTube cluster at the bottom.
- .. **By absolute gap:** The full hierarchy spans 0.65 points; AI sits 0.62 points behind the top source and 0.03 points above the bottom source.

The two readings answer different questions. The report cites the rank–order reading when describing source position, and cites the absolute–gap reading when describing range and band width.

4.4 AI rank by category

AI trust varies by category. The full 18–category breakdown is reported in §5.1. Range: 4.30 (Apparel) to 4.70 (Electronics). Electronics is the only category in which AI's mean trust score crosses the 4.7 mark.

4.5 Broader trust landscape: retailer trust and review trust

Two additional measurements describe the trust landscape around AI shopping but do not sit within the AI–specific information–source hierarchy.

Net retailer trust (PC-7). Respondents rated their agreement with the statement *"Online retailers always have my best interest as a customer in mind"* on a 0–10 scale. The mean score was **6.35 out of 10**. Restated as an NPS–style composite (a standard format used in retailer–trust benchmarks): trust promoters (rated 9 or 10) were **26%** of respondents; trust detractors (rated 0–6) were **51%**. Subtracting the two gives a **Net Retailer Trust Score of –25**. The NPS–style cutoff matches the methodology used in the prior Zappos retailer–trust benchmark (which produced approximately –60 net), allowing the two to be compared on like terms; the methodology rationale is documented at §1.6b.

Review trust (AT-7). Among respondents who reported using online customer reviews in their product research (n=728 from PC-4), confidence in online product reviews was rated on a 1–7 scale ("Not at all confident" to "Extremely confident"). The mean was **4.92 out of 7**. 28% of review users sit in the top–two–box (high confidence). 7% sit in the bottom–two–box (low confidence).

Together, the seven-source information trust hierarchy (§4.1) and these two broader trust measures describe how shoppers' general confidence in retail-information sources compares to their confidence in the retailers and reviews they encounter directly.

(Net retailer trust: n=1,463; source: §4.5; methodology integrity-flag: §1.6b. Review trust mean: n=728 actual review users (from PC-4); see Methodology – Mixed-format encoding correction.)

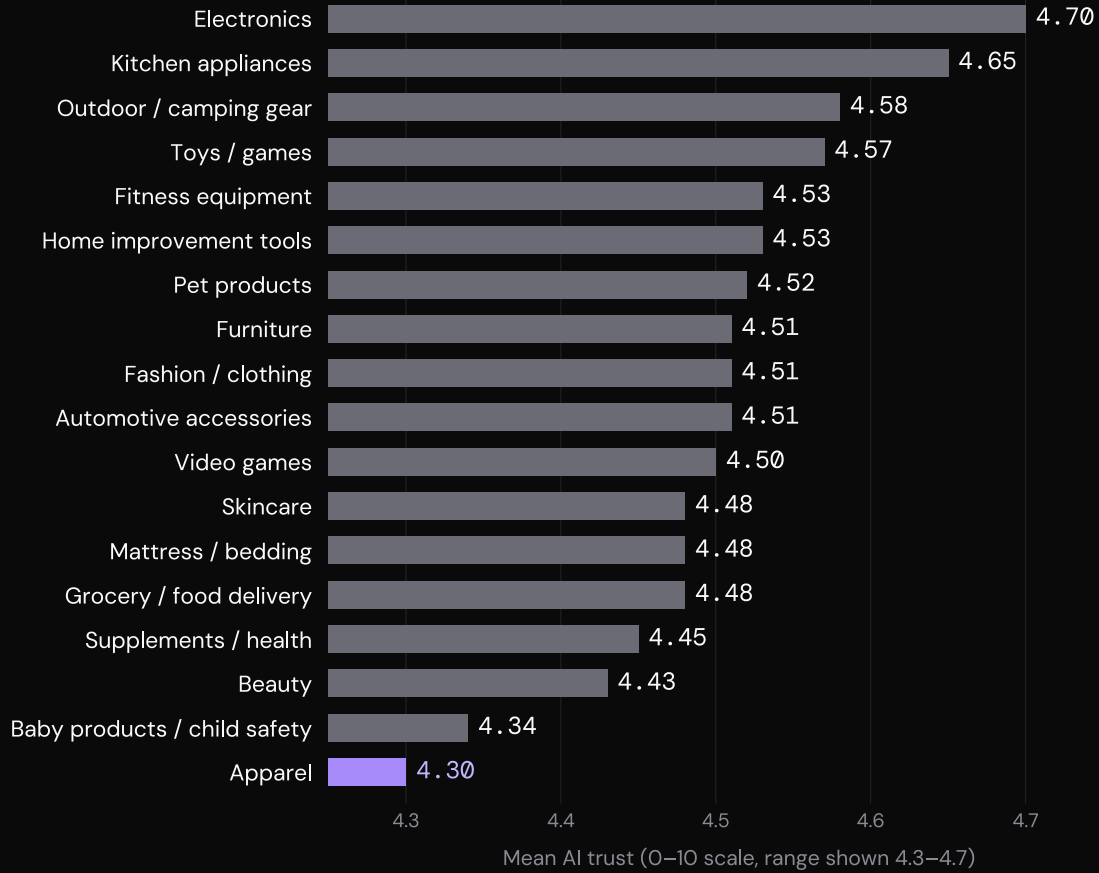
Category & Demographic Patterns

AI shopping behavior varies by category and by demographic. The chapter covers per-category AI trust across the 18 categories surveyed, breakdowns by age and gender, household income, ethnicity, and the strongest category × demographic intersections.

5.1 The Category-Specific AI Trust Index

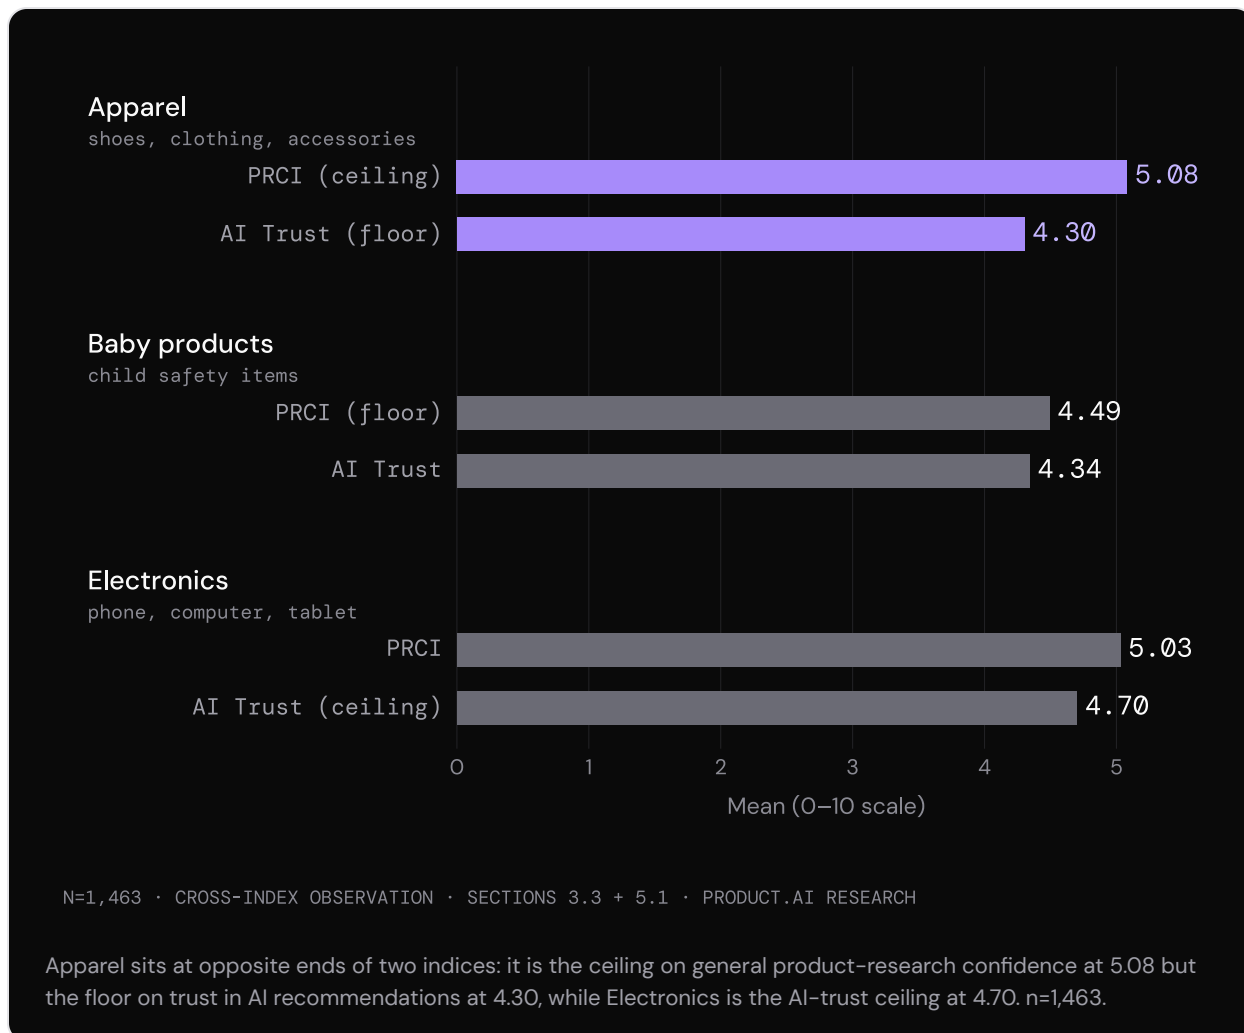
Respondents rated their trust in AI-generated product recommendations across 18 product categories on a 0–10 scale. The category-level means across all 1,463 respondents:

CATEGORY	AI TRUST INDEX (0-10)	TIER
Electronics (smart phone, computer, tablet)	4.70	Highest
Kitchen appliances	4.65	High
Outdoor / camping gear	4.58	High
Toys / games	4.57	Moderate
Fitness equipment	4.53	Moderate
Home improvement tools	4.53	Moderate
Pet products	4.52	Moderate
Furniture	4.51	Moderate
Fashion / clothing	4.51	Moderate
Automotive accessories	4.51	Moderate
Video games	4.50	Moderate
Skincare	4.48	Moderate
Mattress / bedding	4.48	Moderate
Grocery / food delivery	4.48	Moderate
Supplements / health products	4.45	Lower
Beauty (makeup, cosmetics)	4.43	Lower
Baby products / child safety	4.34	Lower
Apparel (shoes, clothing, accessories)	4.30	Lowest



N=1,463 · 0-10 AI TRUST SCALE ACROSS 18 CATEGORIES · RANGE 0.40 POINTS · PRODUCT.AI RESEARCH

Trust in AI-generated product recommendations spans just 0.40 points across 18 categories on a 0-10 scale, from an Electronics ceiling of 4.70 down to an Apparel floor of 4.30. n=1,463.



Electronics tops the index at 4.70. Apparel sits at the bottom at 4.30. The full range spans 0.40 points.

The Apparel cross-index pattern. Apparel is the highest-confidence category in the per-category Purchase Research Confidence Index (5.08 in PC-6, §3.3) and the lowest-trust category in the per-category AI Trust Index (4.30, this section). The same category occupies the ceiling in one measure and the floor in another.

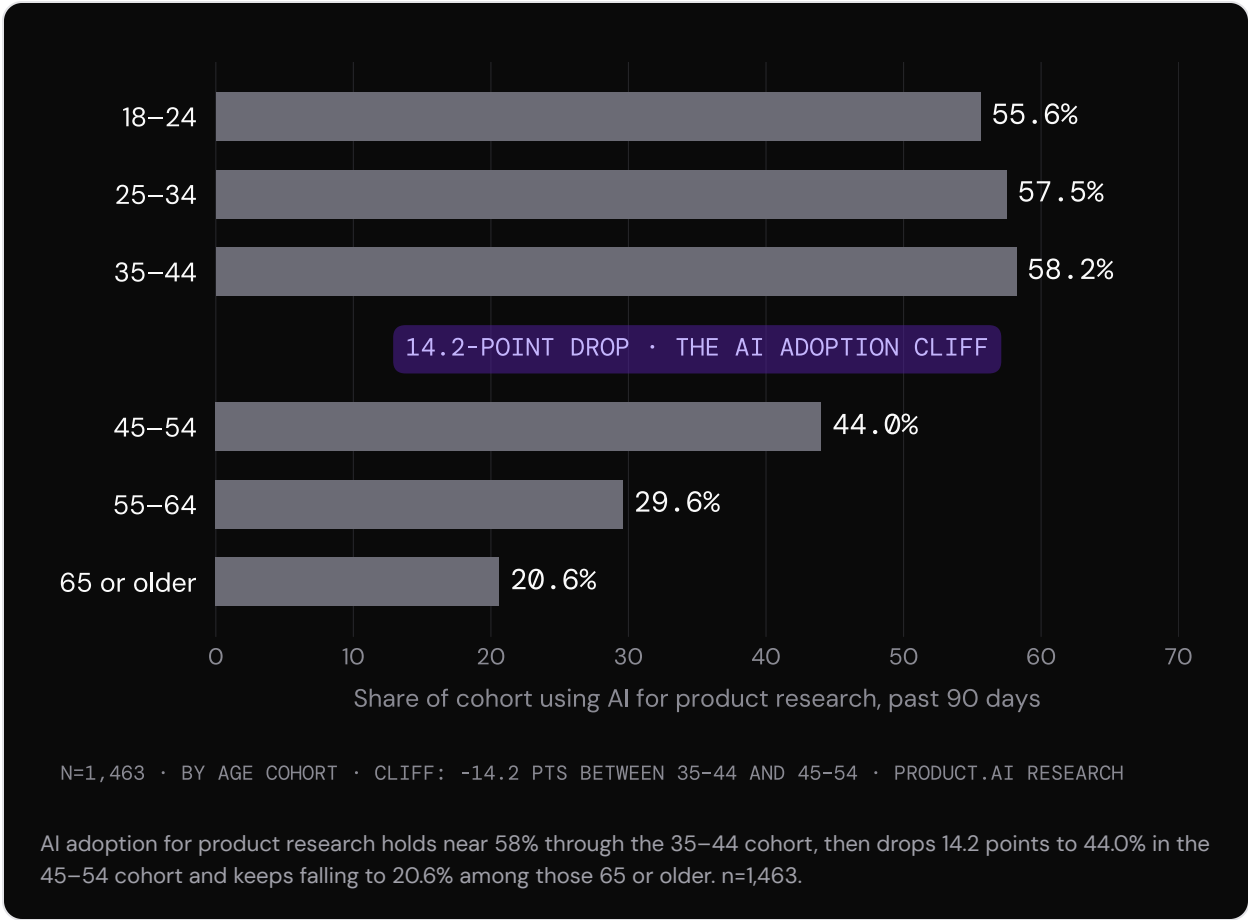
(Category-Specific AI Trust Index: n=1,463; source: §5.1. Cross-index observation: §3.3 + this section.)

5.2 Age and gender

By age

AI usage for product research in the past 90 days varies by age. The full age-bracket distribution across n=1,463:

AGE	N	AI USED PAST 90 DAYS
18-24	171	55.6%
25-34	259	57.5%
35-44	239	58.2%
45-54	232	44.0%
55-64	247	29.6%
65 or older	315	20.6%



The drop between the 35-44 and 45-54 cohorts is 14.2 points. The drop between 45-54 and 55-64 is 14.4 points. Under-44 cohorts cluster between 55.6% and 58.2% AI usage. 55+ cohorts

measure below 30%.

By gender

MEASURE	MALE (N=737)	FEMALE (N=726)
AI used past 90 days	50.1%	35.0%
AI trust source rating mean (0–10)	4.69	4.30
Verification rate among AI users	88.9% (n=369 AI users)	81.1% (n=254 AI users)
Autonomy cap under \$25	33.7%	50.0%

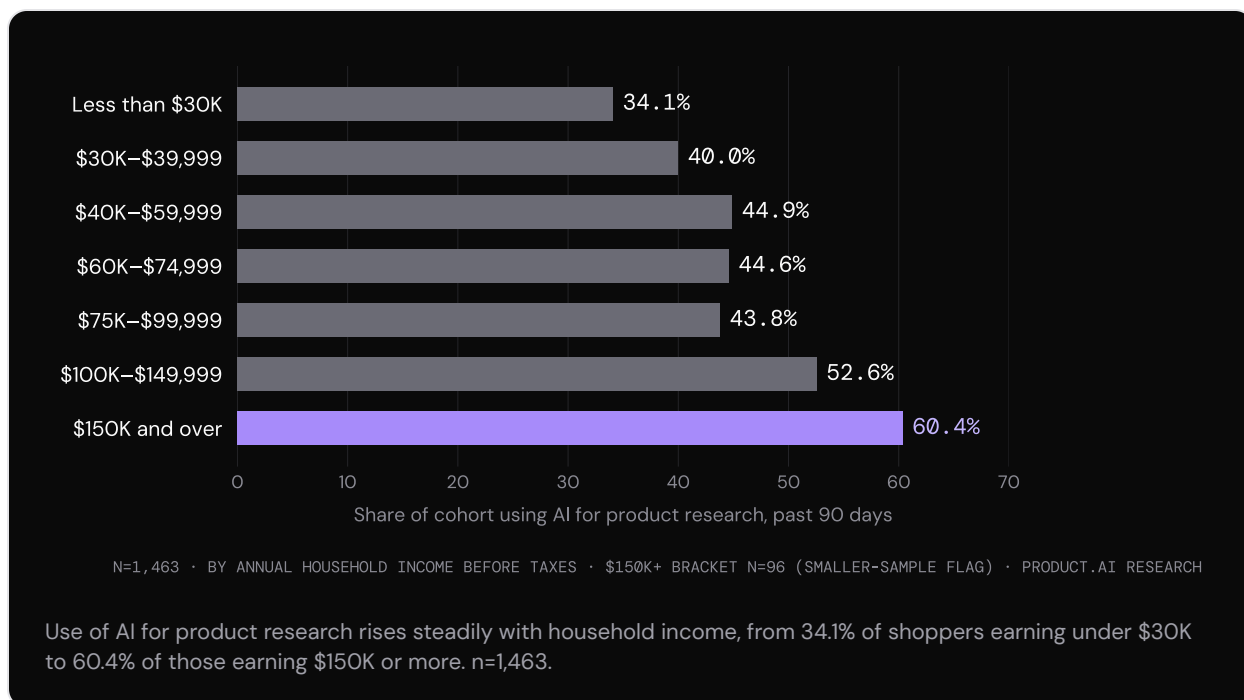
Male respondents reported AI usage at 1.43× the female rate (50.1% vs 35.0%, a 15.1–point gap). Female respondents more often cap their AI trust under \$25 (50.0% vs 33.7%, a 16.3–point gap). Verification rate among AI users runs 7.8 points higher for male AI users (88.9% vs 81.1%).

(Age and gender cuts: source data n=1,463; see §5.2.)

5.3 Household income

AI behavior varies across the seven income brackets surveyed. Two brackets are flagged for smaller sample size: \$60K–\$74K (n=148) and \$150K+ (n=96). Findings at those brackets should be read with the smaller sample in mind.

INCOME	N	AI PAST 90D	VERIFICATION (AI USERS)	AUTONOMY <\$25	LAST-PURCHASE CONFIDENCE (0–10)
Less than \$30K	443	34.1%	80.1%	55.7%	5.74
\$30K–\$39,999	205	40.0%	76.8%	36.1%	5.56
\$40K–\$59,999	243	44.9%	84.4%	37.4%	6.21
\$60K–\$74,999 △	148	44.6%	89.4%	37.8%	6.59
\$75K–\$99,999	176	43.8%	92.2%	35.2%	6.94
\$100K–\$149,999	152	52.6%	93.8%	38.2%	7.30
\$150K and over △	96	60.4%	91.4%	25.0%	7.28



Across the seven brackets:

- AI adoption rises from 34.1% (lowest bracket) to 60.4% (highest bracket) — a 26.3-point spread.
- Verification rate among AI users rises from 76.8% in the \$30K–\$39,999 bracket to 93.8% in the \$100K–\$149,999 bracket — a 17-point spread.
- The share of respondents capping AI trust under \$25 falls from 55.7% in the lowest bracket to 25.0% in the highest — a 30.7-point spread.
- Last-purchase confidence rises from 5.74 to 7.28 on the 0–10 scale — a 1.54-point spread.

(Income breakdowns: source data n=1,463; see §5.3.)

5.4 Ethnicity

Respondents identified their ethnicity. The achieved sample distribution: 759 White (non-Hispanic), 285 Hispanic or Latino, 186 Black or African American, 147 Asian or Pacific Islander, 86 Other. The "Other" cell is below the 100-respondent threshold and is flagged accordingly.

ETHNICITY	N	AI PAST 90D	VERIFICATION (AI USERS)	AUTONOMY <\$25	AI TRUST SOURCE RATING (0-10)
White (non-Hispanic)	759	37.4%	87.3%	48.5%	4.52
Hispanic or Latino	285	49.5%	86.5%	29.5%	4.55
Black or African American	186	50.0%	76.3%	29.0%	4.55
Asian or Pacific Islander	147	53.1%	87.2%	34.7%	4.53
Other Δ	86	31.4%	92.6% (n=27 AI users)	62.8%	4.09

Across the four larger groups:

- AI adoption is 12.1 to 15.7 points higher among Hispanic, Black, and Asian/PI respondents than among White respondents (49.5%, 50.0%, 53.1% versus 37.4%).
- Verification rate among AI users runs 76.3% for Black AI users; the other three larger groups cluster between 86.5% and 87.3%.
- White respondents most commonly cap AI trust under \$25 (48.5%). Hispanic and Black respondents do so least often (29.5%, 29.0%).
- The AI trust source rating means cluster tightly across the four larger groups (4.52–4.55).

(Ethnicity breakdowns: source data n=1,463; "Other" cell n=86 / AI users within Other n=27; see §5.4.)

5.5 Category × demographic intersections

The strongest within-category demographic gaps appear in two cuts: category × gender and category × age. The values are AI Trust Index means on the 0–10 scale.

Category × gender — largest male–female gaps

CATEGORY	MALE	FEMALE	M-F GAP
Automotive accessories	4.70	4.26	+0.44
Mattress / bedding	4.67	4.25	+0.42
Fitness equipment	4.71	4.31	+0.40
Furniture	4.69	4.30	+0.39
Electronics	4.87	4.49	+0.38

Across all 18 categories, the male mean is at or above the female mean. The category-level gap ranges from +0.44 (Automotive) to +0.06 (Baby Products).

Category x age — largest under-44 vs 55+ gaps

CATEGORY	UNDER 44 (N=669)	55+ (N=562)	GAP
Beauty	4.61	4.11	+0.50
Baby products	4.50	4.01	+0.49
Fashion / clothing	4.66	4.27	+0.39
Grocery / food delivery	4.61	4.24	+0.37
Supplements	4.55	4.19	+0.36

Mapping the two cuts to each other

Categories with the largest gender gaps (Automotive, Mattress, Fitness, Furniture, Electronics) sit in the bottom half of the age-gap ranking. Categories with the largest age gaps (Beauty, Baby products, Fashion, Grocery, Supplements) sit in the bottom half of the gender-gap ranking. The two demographic-cut patterns map to different category sets within the 18-category space — categories where gender splits AI trust most are not the same categories where age splits AI trust most.

(Intersections: source data n=1,463 cross-tabbed by gender (n=737 Male / n=726 Female) and by age cohort (n=669 under-44 / n=562 55+); full 18-category decompositions documented in the demographic crosstab extract.)

Open Research Questions and Planned Additions

The research is intended as the inaugural round of a recurring program. Four questions emerged that future rounds can answer with comparable measurements. Four instrument additions are planned for future rounds.

6.1 Instrument additions planned for future rounds

Four questions or measures were either deferred from this round or identified as needed during analysis. Each enters the next round of the research:

- **A dedicated verification-source-stack question.** Section 2.3 reports that 86% of AI users verify the AI's recommendation through another source. Section 2.4 reports which tools shoppers use for product research broadly. The research did not separately measure which

specific sources AI users consult when they verify an AI recommendation. A dedicated question on the verification–source stack is the highest–priority instrument addition.

- .. **Savings–tool preference (SB–3).** Measurement of relative preference among savings tools was deferred from this round to keep the instrument compact and to allow calibration against the base measurements collected in this round. It enters the next round.
- .. **Question–level structured open–text capture.** This round used a single open–text closer. Future rounds add structured open–text prompts at the highest–information points (post–purchase reflection, reasoning behind blame attribution, reasoning behind retailer–intent beliefs, what would have made the respondent more confident in the purchase).
- .. **Strengthened skip logic at SB–1.** A skip–logic gap in this round routed 271 respondents who reported rarely or never looking for promo codes into the full Checkout Gap block. Future rounds bypass the Checkout Gap block for those respondents, tightening the denominator for SC–side stats reported in the companion industry report. The gap is documented in Methodology – Integrity flags and limitations §7.8.

6.2 Open research questions

Four research questions are open at the close of this round. Each is structured as a measurement that future rounds will produce a comparable value for.

- .. **Does the 86% verification rate hold across rounds?** This round measures verification at 86% among AI users (n=623). Future rounds re–measure the same question on the same denominator definition.
- .. **Is the \$25 AI Autonomy Threshold structural?** This round measures 42% capping AI trust under \$25, with the full dollar–tier curve in §2.2. Future rounds re–measure the same six brackets. Whether the under–\$25 share rises, falls, or holds is a direct comparison.
- .. **Does the per–category Purchase Research Confidence Index ceiling or floor shift?** This round measures the PC–6 ceiling at 5.08 (Apparel) and the floor at 4.49 (Baby Products), with the full 18–category table in §3.3. Future rounds re–measure all 18 categories on the same scale.
- .. **Does the AI position in the information trust hierarchy move?** This round measures AI at #6 of 7 at 4.51, with the top source at 5.13 and no source above 5.2. Future rounds re–measure the same seven sources.

6.3 Cadence and comparability

The instrument is structured for replication. Question wording, response scales, response brackets, and category lists are held constant across rounds. The same sample panel (Cint general-population, U.S. online shoppers) is used.

Future rounds will introduce Census-weighted post-stratification on age, gender, region, and income — described in Methodology — Weighting — to allow like-for-like comparison across rounds. The first inter-round comparison will re-weight this round's data to the same post-stratification profile so the two reads are directly comparable.

6.4 What this round establishes as baseline

The values reported in this round serve as the baseline for each of the named indices and measures the research tracks:

- **Purchase Research Confidence Index (PRCI).** PC-1 last-purchase mean: 6.95 out of 10 (n=1,453). PC-6 per-category range: 4.49 (Baby Products) to 5.08 (Apparel), cross-category mean approximately 4.98.
- **AI Autonomy Threshold.** 42% cap AI trust under \$25. Full dollar-tier distribution: §2.2.
- **Verification rate among AI users.** 86% verify the AI's recommendation through another source.
- **Information trust hierarchy.** AI at #6 of 7 at 4.51; top source 5.13; range 0.65 points across all seven sources.
- **Category-Specific AI Trust Index.** Range 4.30 to 4.70 across 18 categories. Full table: §5.1.
- **Net Retailer Trust Score (PC-7).** -25 NPS-style composite; mean 6.35 out of 10.

Each is a single quantifiable value or a structured table that the next round re-measures on identical question structure.

Limitations

Reader's note. The findings in this report describe what the research measured. The limitations describe what the research did not measure, what reads cautiously, and what future research will deepen. The two layers belong together. A reader who wants to audit any specific limitation can follow its sub-section into the methodology chapter (§1) for the underlying instrument design or sample-frame definition.

7.1 Point-in-time snapshot

This research is a point-in-time measurement, fielded on April 27, 2026. It captures a snapshot of U.S. online-shopper behavior, trust, and verification practice in Q2 2026. Trend-level claims wait for cross-round deltas in future research.

Specifically: a reader should not interpret any reported value as a directional trend. The 86% verification rate is a Q2 2026 baseline. The 42% AI Autonomy Threshold is a Q2 2026 baseline. The Purchase Research Confidence Index value of 6.95/10 is a Q2 2026 baseline. Each will be re-baselined in future research; only then do cross-round deltas describe a trend.

The point-in-time constraint also limits causal interpretation. The research observes association, not causation. Where the report describes patterns (e.g., the relationship between purchase price tier and AI autonomy), the patterns are descriptive correlations measured at a single point in time, not causal claims tested across time.

7.2 U.S.-only sample frame

The sample is drawn exclusively from U.S. residents. International generalization is not supported by the data. A reader who wants to apply findings to non-U.S. consumer behavior should consider those findings hypothesis-generating, not confirmatory.

The U.S. sample frame was a design choice, not a sampling artifact. The instrument's screener excluded non-U.S. respondents at intake. Future research retains the U.S. frame; expansion to additional national markets is a candidate question for a subsequent round, not a current commitment.

7.3 AI-user subsample for verification stats

The 86% verification rate — the report's headline canonical statistic — uses the AI-user subsample (n = 623 of 1,463 total). Statistical inference applies to the AI-user population (U.S. online shoppers who used AI for product research in the past 90 days), not to the broader online-shopper population.

This is a deliberate methodological choice, documented in advance in the Research Brief Part 13. The verification behavior is only meaningfully measurable among respondents who actually used AI for product research; applying the full-sample denominator would understate the rate by mixing AI users with non-users who could not have verified an AI recommendation. The methodology chapter sub-section §1.5b documents the denominator rationale in full.

A reader citing the 86% statistic should report the denominator alongside the rate. The footnote and citation convention (Methodology — Footnote and citation convention) provides the

canonical citation form.

7.4 Online shoppers only (not all U.S. consumers)

The sample frame is defined as U.S. adults who shop online. Offline-only consumers — adults who do not shop online, or shop online rarely enough to fail the screener qualification — are excluded by design.

Findings generalize to the U.S. online-shopper population, not to all U.S. consumers. Estimates of total-population AI usage, total-population trust in retailers, or total-population promo-code behavior are not supported by this research; the frame excludes the relevant denominator.

This constraint matters most for any reader applying these findings to total-population marketing planning. The online-shopper subset over-indexes on younger and higher-income demographics relative to the total U.S. adult population.

7.5 18 categories surveyed — not exhaustive

The research measures the Purchase Research Confidence Index and the AI Trust Index across 18 product categories. The categories represent the most-shopped consumer categories surveyed against; less-shopped categories (e.g., heavy machinery, specialty B2B goods, professional services) are not in scope.

Generalization of either index beyond the surveyed 18 categories is not supported. A reader should not project the 4.30–4.70 AI Trust Index range onto unmeasured categories without independent measurement.

Future research retains the same 18-category structure for like-for-like cross-round comparison. Category expansion is a candidate question for a subsequent round.

7.6 Self-reported behavior

All behavioral measures are self-reported. Respondents describe their own behavior; the instrument does not observe behavior independently.

Social-desirability bias affects self-report directionally on socially-shaped behaviors. The 86% verification rate is the load-bearing case: respondents may over-report verification because verification is the socially-aligned answer. The directional finding — verification is the dominant AI shopping behavior — holds robustly across the instrument; the precise rate carries self-report uncertainty.

Other self-report-affected measures include the post-purchase regret statistics (15% wished they had researched more; 13% found a better option afterward), the AI satisfaction score, and

the research-time measure. The report cites each as self-reported; readers should treat the directional reading as more reliable than the precise percentage.

7.7 Small-cell footnotes (where they apply)

Cross-tabs of the sample maintain publishable cell sizes (n=80+) at the Age × Gender level. The demographic crosstab work documents the cell-size discipline.

Sub-cuts that approach n<30 — for example, Age × Gender × Category three-way intersections in less-represented combinations — are footnoted in the report body where they appear. A reader encountering a small-cell footnote should treat the specific sub-cut as directional only.

Section 6 (Category & Demographic Patterns) carries the small-cell disclosures relevant to the report's demographic findings — including the n=148 and n=96 income brackets in §5.3 and the n=86 "Other" ethnicity cell (n=27 AI users within) in §5.4.

7.8 What is planned for future research

The research was instrumented as the inaugural round of a recurring program, not as a one-off study. Three structural elements were deliberately held back for future rounds:

- .. **Savings-tool preference (SB-3).** Measurement of relative preference among savings tools was deferred to future research to keep this round's instrument compact and to allow calibration against the base measurements collected here.
- .. **Structured open-text capture.** Free-text response capture for selected questions was held back to future research; this round used a single Open-Text Closer block at instrument end. Future research adds question-level open-text prompts at the highest-information points.
- .. **Strengthened SB-1 skip logic.** A skip-logic gap in this round (documented as integrity flag #1.6 in the methodology chapter) routed respondents who reported rarely or never looking for promo codes into the full Checkout Gap block. The gap affects SC-side stats reported by code-user denominator; PAI-side Trust Report stats are not affected by this gap. Future research resolves the gap by bypassing the CG block for rarely-or-never respondents.

Future research also introduces Census-weighted post-stratification (see Methodology — Weighting) for like-for-like cross-round comparison.

7.9 Mixed-format encoding correction (pre-publish)

During pre-publish review, an exclusion artifact was identified in the original analysis of three scale questions and corrected before publication. The raw Alchemer export mixed text labels (e.g., "Extremely confident," "Neutral," "Not at all confident") and numeric values on the same

scale column for endpoint-anchored items. Initial mean computations dropped the text-labeled responses, which silently understated three published means by excluding the highest-frequency endpoint clusters from the average.

Three values are restated against the corrected calculation, with the original published-stat-value and the corrected-stat-value disclosed alongside:

STAT	ORIGINAL (MIXED-FORMAT-EXCLUSION)	CORRECTED (FULL-MAPPING)
PRCI Last-Purchase Confidence mean (§3.2)	6.34 / 10 (n=869 pure-numeric responders, mislabeled as n=1,463)	6.95 / 10 (n=1,453 full non-missing)
AI Satisfaction mean (§1.3)	5.06 / 7 (n=309 pure-numeric, mislabeled as n=623)	5.46 / 7 (n=623 AI users)
Review-trust denominator (§4.5)	n=820 "review users"	n=728 actual review users (PC-4 selected "Online customer reviews"); mean 4.92 / 7

The systematic direction of the original error: means understated, because the "Extremely [positive]" endpoint cluster of each scale was the largest single response group. No other canonical stat in the report is affected; per-category PRCI and AI Trust Index rankings, the Information Trust Hierarchy means, the AI Autonomy Threshold distribution, AI usage rates, AI verification rates, retailer-trust scores, and all demographic breakdowns are accurate against the raw data as originally computed.

The remap rule applied: text-labeled endpoint responses are mapped to the numeric anchor each label represents — for PC-1 ("Extremely confident" → 10, "Neutral" → 5, "Not at all confident" → 0 on the 0-10 scale); for AT-3 and AT-7 ("Extremely [X]" → 7, "Neutral" → 4, "Not at all [X]" → 1 on the 1-7 scale) — prior to averaging across the full non-missing column. The rule is documented in the [companion codebook](#) under "Mixed-format encoding correction" so anyone replicating the analysis from the raw dataset applies the same mapping. Disclosed here rather than silently corrected because the report's thesis is that verifiable evidence requires methodology transparency.

Methodology

Sample frame rationale

The Trust in AI Commerce Report measures consumer behavior among U.S. online shoppers. The sample frame is defined by three qualifying conditions: (1) U.S. residency, (2) age 18 or older, and

(3) self-reported online shopping in the recent past. The screener excluded non-U.S. residents, minors, and respondents who reported no online shopping activity.

Online shoppers — as distinct from all U.S. consumers — were chosen as the sample frame for two reasons. First, the report measures behavior at the surface where AI product recommendations are most actively encountered: online retail. Offline-only shoppers do not interact with AI shopping assistants at the rate online shoppers do, and including them would dilute the behavioral signal. Second, the instrument captures behavior across both e-commerce trust patterns (the Checkout Gap block) and AI-shopping patterns (the AI Trust block); both measurement layers require respondents with active online-purchase context.

The frame definition has methodological implications disclosed throughout this chapter. Estimates produced from the dataset generalize to the U.S. online-shopper population, not to all U.S. consumers. Findings on AI-specific behaviors apply to a further sub-population: respondents who reported using AI for product research in the past 90 days (n = 623). The denominator for each canonical stat in this report is stated explicitly at the point of citation.

Fielding details

PARAMETER	DETAIL
Field date	April 27, 2026
Survey platform	Alchemer
Sample panel	Cint
Sample frame	U.S. adults who shop online
Total records collected	3,638
Disqualified at screen	1,778
Partial completions	397
Complete responses	1,463
Analysis basis	Complete responses only (Status = "Complete")
Rounding	Percentages rounded to one decimal place; mean scores rounded to two decimal places

The survey was a cross-sectional snapshot fielded on April 27, 2026 through the Cint general-population panel and administered through the Alchemer platform. Field completion was rapid because the panel was pre-recruited; no longitudinal cohort tracking is implied by the design.

Of 3,638 panelists who began the instrument, 1,778 were disqualified at the screen (primarily for failing the U.S. residency or online-shopping eligibility checks), 397 left the instrument before completion, and 1,463 produced complete responses that passed the quality controls described in §1.7. Only the 1,463 complete responses appear in the analysis dataset.

Screener logic

Three screener questions preceded the main instrument:

- .. **U.S. residency.** Respondents outside the United States were disqualified.
- .. **Age threshold.** Respondents under 18 were disqualified.
- .. **Online-shopping qualification.** Respondents who reported no online shopping activity were disqualified.

Within the qualified sample, two additional sub-population qualifications routed the instrument:

- Respondents who reported using promo codes in the past 30–60 days (n = 876, or 59.9% of the qualified sample) received the full Checkout Gap (CG) block. Respondents who did not use promo codes in that window skipped the CG block; the SC-side canonical stats anchored on this n=876 sub-sample, not the full n=1,463.
- Respondents who reported using AI for product research in the past 90 days (n = 623, or 43% of the qualified sample) received the AI verification-behavior questions. The PAI-side 86% verification stat anchored on this n=623 sub-sample.

The full instrument is published alongside this report as 2026-04-Checkout_Gap_Gen_Pop_Survey.docx.

Sample composition

The achieved sample (n = 1,463) shows the following demographic distribution. Cint panel quotas balance age, gender, and region; sample composition was not weighted to U.S. Census distributions in this round. Where sample composition deviates from U.S. online-shopper population estimates, the deviation is disclosed in §1.8.

AGE GROUP	N	SHARE
18–24	171	12%
25–34	259	18%
35–44	239	16%
45–54	232	16%
55–64	247	17%
65+	315	22%

GENDER	N	SHARE
Male	737	50%
Female	726	50%

ANNUAL HOUSEHOLD INCOME	N	SHARE
Less than \$30,000	443	30%
\$30,000 – \$39,999	205	14%
\$40,000 – \$59,999	243	17%
\$60,000 – \$74,999	148	10%
\$75,000 – \$99,999	176	12%
\$100,000 – \$149,999	152	10%
\$150,000+	96	7%

Gender balance is exact (50/50). Age distribution skews moderately older – the 65+ bracket holds 22% of the sample, the largest single age group. Income distribution skews lower than the U.S. online-shopper mean, with 30% of respondents reporting household income under \$30,000.

Denominator definitions

Every canonical stat in this report is paired with the denominator from which it was computed. The following sub-anchors define the load-bearing denominators used across the report. A reader who needs to verify a claim should follow the in-line denominator citation to its sub-section here.

1.5a Full sample (n = 1,463)

The full analysis dataset of complete responses. Used for all canonical stats that apply to the U.S. online-shopper population at large, including: the AI Autonomy Threshold (42% of all

respondents won't trust AI for purchases over \$25 without verifying first), AI usage rates (43% used AI for product research in the past 90 days; 20% used AI for their most recent \$50+ purchase), the Purchase Research Confidence Index overall measure (mean 6.95/10), and the information trust hierarchy ranking (7 sources with means 4.48–5.13).

1.5b AI-user subsample (n = 623)

The subset of complete respondents who reported using AI for product research in the past 90 days. Used for the AI verification-behavior stats: 86% of AI users verified the AI's recommendation through another source before buying (45% always verify + 41% sometimes verify = 86%). Only 14% trusted the AI recommendation without verification (n = 89 of 623).

The AI-user denominator is the correct base for verification-behavior claims because respondents who did not use AI for product research could not have verified an AI recommendation. Applying the full-sample denominator to the verification stat would understate the verification rate by mixing AI-users with non-users. The choice of denominator was documented in advance in the Research Brief (Part 13) and is consistent with established sub-population reporting discipline.

1.5c Code-user subsample (n = 876)

The subset of complete respondents who reported using a promo code in the past 60 days. This denominator anchors the SC-side Checkout Verification Index stats (35% experienced code failure; 13% experienced full failure; 22% experienced partial failure). PAI-side Trust Report findings do not anchor on this denominator; it is documented here for cross-report integrity.

1.5d Other sub-sample denominators

Selected questions routed through skip logic to narrower sub-populations. Where a Trust Report finding cites a sub-sample denominator other than the three above (for example, the n=728 review-users denominator (from PC-4) for the AT-7 review-trust mean), the denominator and qualifying condition are stated at the point of citation.

Integrity flags

Two integrity considerations affecting the canonical-stat baselines were surfaced and resolved during analysis. Both are disclosed here so a reader can audit the resolution.

1.6a AI Autonomy Threshold – rounding to 42%

The 42% AI Autonomy Threshold reflects the share of respondents whose maximum spend on an AI recommendation without checking another source falls under \$25. The unrounded value is

41.8% (n = 611 of 1,463). The report cites 42% (rounded to whole-percent for headline clarity); the unrounded 41.8% appears in the data table at §2.

1.6b PC-7 Net Trust Score — NPS-style cutoff

The PC-7 question measured retailer trust on the statement "Online retailers always have my best interest as a customer in mind" using a 0–10 scale. The Net Trust Score (–25) was computed using the standard Net Promoter Score (NPS) cutoff convention: trust promoters (scores of 9–10) at 26%, minus trust detractors (scores of 0–6) at 51%, equals –25. The NPS-style cutoff is the methodology used in the prior Zappos retailer-trust benchmark (~–60 net) referenced in the research design; this round's finding diverges from that prior benchmark in direction (less negative) but holds the same cutoff convention for comparability.

The mean score on the same question is 6.35/10. The report cites both the mean (6.35) and the Net Trust Score (–25) where each is most informative; readers should not interpret the –25 as a directional rate. It is an NPS-style composite, not a percentage.

Quality controls

Four quality-control measures were applied during fielding and analysis:

- .. **Screener disqualification.** Respondents failing the U.S. residency, age, or online-shopping screeners were disqualified at intake (n = 1,778 disqualified at screen). This protects the sample frame.
- .. **Completion gate.** Only respondents who completed the full instrument (Status = "Complete") entered the analysis dataset. Partial completions (n = 397) were excluded.
- .. **Skip logic enforcement.** Sub-sample routing — to the Checkout Gap block (code users) and the AI verification block (AI users) — was enforced at instrument level, preventing respondents from answering sub-population questions outside their qualifying condition. A skip-logic gap at SB-1 ("rarely or never look for promo codes" respondents were still routed into the full CG block) is documented in the limitations chapter (§8) and resolved for future rounds.
- .. **Attention and speed.** The instrument included pacing checks. Respondents whose completion times fell below the panel-base speeder threshold were filtered before the dataset entered analysis. The exact threshold is documented in the Alchemer field report.

Weighting

The survey was not weighted to U.S. Census distributions. Cint panel quotas balanced age, gender, and region at the panel level; further post-stratification weighting was not applied. The

achieved sample (n = 1,463) skews moderately older and lower-income than U.S. online-shopper population estimates. Where this skew affects a finding's generalization to the U.S. online-shopper population, the limitations chapter (§8) discloses the directional risk.

Future research will introduce Census-weighted post-stratification on age, gender, region, and income to improve generalization precision. Cross-round deltas reported in subsequent research will compare like-for-like by re-weighting this round's data to the same post-stratification profile.

Limitations summary

A short callout: *This research is a U.S.-only, self-reported, point-in-time snapshot of Q2 2026 behavior. The 86% verification stat applies to AI users (n = 623), not all online shoppers. Categories surveyed are 18 of the most-shopped; less-shopped categories are not represented. Self-report bias affects directional reading on socially-shaped behaviors (verification, post-purchase regret). The full structured limitations block is at §7.*

Footnote and citation convention

Every canonical stat in this report carries an in-line footnote in the following format:

(value%, denominator description, source: §anchor)

Worked example for the 86% AI verification stat as cited in §2:

86% of U.S. online shoppers who use AI for product research verify the AI's recommendation through another source before buying. (n = 623 AI users among 1,463 surveyed; 45% always verify + 41% sometimes verify; source: §2.3)

This footnote convention is consistent across all canonical stats. A reader who follows the source link from any in-line citation lands on the sub-section of this chapter where the relevant denominator and integrity considerations are defined. This is the substrate that makes every stat in the report independently auditable.
