

Consumer AI Trust Is Task Specific

An executive guide to deciding what consumers will let an assistant do

Consumer AI Adoption Gap · Research brief 03
October 3, 2026

© 2026 Eldris Inc. All rights reserved.

Executive brief

A consumer uses an AI assistant to compare products. Does that person also want the assistant to choose a product, access payment details and complete the purchase?

Those are separate decisions. A general measure of AI trust cannot tell a product team where one permission ends and the next begins. Nor does a successful recommendation establish that the consumer will accept recurring authority, a substitution or the handling of an unexpected problem.

The commercial question is specific: **for which task, under which conditions, will a consumer grant permission and continue to rely on the service?**

This report uses shopping as a concrete setting for that question. Its intended readers are consumer AI product leaders, retailers and payment providers deciding whether to offer assistance, require transaction approval or test bounded delegation. Other tasks can use the same decision framework, but shopping findings should not be transferred directly to health, financial advice or personal relationships.

The main recommendation is to make **a defined permission**, rather than generalized trust, the unit of product research. Specify the task, action, data, limits and recovery process. Then measure what people actually authorize, whether they understand it, and what happens over subsequent uses.

Three requirements follow:

- **Separate assistance from action.** Search, recommendation, purchase approval and continuing authority need their own evidence.
- **Make control usable.** People should understand what the assistant may do, change its scope and withdraw future authority.
- **Evaluate continued reliance.** Initial permission is only part of the outcome; corrections, regret, revocation and repeat use matter too.

The objective is appropriate reliance on a service that delivers value within understood boundaries. More delegation is not automatically a better product outcome.

What public evidence can tell us

Public research helps locate questions to investigate. It provides less certainty about which authorization design will work in a particular product.

Pew Research Center's February 2026 survey of 5,119 U.S. adults found that 49% reported using AI chatbots. That establishes reported exposure and use, not permission to spend money or confidence in a particular task. [Pew report and methodology](#).


Gartner's January 2026 survey of 322 U.S. consumers found 31% willing to let AI narrow household-supply choices and 28% willing to let it narrow personal-electronics choices. Willingness to let AI make purchase decisions reached at most 11% across lower-stakes categories. These are reported intentions, not transaction observations; the ceiling is not a category-matched conversion rate. [Gartner release](#).

The distinction becomes more detailed in ACI Worldwide's YouGov survey of 2,080 UK adults, fielded June 19-22, 2026. Reported trust ranged from 50% for finding the best price to 15% for handling problems. The responses concern different capabilities; they are not measured product success rates. [ACI research announcement](#).

CONSUMER AI ADOPTION GAP / RESEARCH BRIEF 03

Trust differs across shopping capabilities

Share of UK adults reporting trust in AI for each capability



Stated trust, not observed delegation or a test of system performance.

Source: ACI Worldwide / YouGov. UK adults; n = 2,080; fieldwork June 19–22, 2026.
Online, population-weighted survey. Responses may overlap. No significance claim is implied.

© Eldris Inc

Figure 1. Trust differs across shopping capabilities. Shares of UK adults reporting trust, as summarized by ACI. Responses may overlap. The bars are not a conversion funnel or a measure of actual system performance.

The practical implication is that a recommendation feature and a purchase agent should not share one undifferentiated adoption forecast. Research should test the authority the proposed product actually requests.

These surveys do not prove that hesitation is caused by lack of control, poor accuracy, unfamiliarity or another single factor. Treat those as competing explanations to investigate. A person may reject a service because it is unnecessary or costly, even while believing it would work.

Define the task before defining the market


“Shopping” contains several distinct decisions. Reordering a known household item differs from choosing a gift, comparing electronics or entering a financial commitment. Price is only one dimension: substitution tolerance, personal significance, data access, urgency and the consequences of error also matter.

Checkout.com’s June 2026 six-market release reports willingness to delegate grocery shopping at 41%, household supplies at 31% and financial services at 15%. These stated preferences support examining task differences. They do not establish a universal launch order or a U.S. market estimate. [Checkout.com release](#).

CONSUMER AI ADOPTION GAP / RESEARCH BRIEF 03

Delegation preferences vary by category

Reported willingness to delegate shopping · Selected six-market results



Survey preferences, not purchases or a forecast of market conversion.

Source: Checkout.com research release, June 9, 2026. Six consumer markets; not a US-only estimate.

Selected categories. Complete sample and weighting details are not supplied in the cited release.

© Eldris Inc

Figure 2. Delegation preferences vary by category. Selected categories reported in one release; not an exhaustive ranking. The release does not provide sufficient methodological detail to use these figures as precise market forecasts.

For a proposed routine-purchase service, define the research population around relevant behavior: adults who actually buy the selected item category, at the relevant frequency, through the channels the product supports. Decide whether the business question concerns existing assistant users, potential new users or both.

A workable decision might be: should the product offer automatic replenishment of a specified household item, subject to a consumer-set spending limit and substitution rule, or retain approval for every order?

That question has a task, population, competing product designs and a decision. “Do consumers trust AI?” has none of those boundaries.

Segment first around plausible differences in the task: repeat versus unfamiliar purchases, flexible versus strict preferences, and consequences of a wrong outcome. Demographics can add context, but should not substitute for an explanation of the relevant behavior. Report uncertainty when subgroup samples are small.

Product implication: Choose a bounded task and test its authorization conditions. Category-level interest is a reason to investigate, not sufficient evidence to launch autonomous purchasing.

Separate the permissions an assistant requests

An assistant can help a consumer while leaving the final decision entirely with that person. It can also perform a selected action after approval, or operate repeatedly within pre-agreed limits. These arrangements make different demands on the consumer and the product.



Figure 3. Assistance and delegation require different evidence. An original Eldris framework. These are alternative authority arrangements, not a mandatory adoption ladder or a validated prediction of consumer behavior.

Advise. The assistant supplies information or comparisons. The consumer chooses whether to rely on it and completes the action elsewhere. Measure usefulness, accuracy, verification effort and whether the assistance improves the decision.

Prepare. The assistant assembles a proposed cart or transaction. It has permission to prepare, not to buy. Measure whether the proposal matches stated preferences and how much work the consumer needs to correct it.

Execute an approved action. The consumer reviews the material terms and authorizes a particular purchase. Measure comprehension, deliberate approval, fulfillment and departures from the approved terms.

Act within standing limits. The consumer authorizes repeated activity under specified conditions. Measure understanding of the continuing scope, adherence to limits, exceptions, repeat use and withdrawal of permission.

These arrangements need not represent increasing maturity. A product may remain valuable indefinitely as an adviser. A consumer can prefer transaction approval for one category and standing permission for another.

The permissions should also specify relevant data use. Access to a shopping history, an address or a payment method changes the arrangement. Permission for a recommendation should not silently become permission to connect accounts, make payments or broaden data use.

For standing permission, a short authorization record should answer: what may happen, for which items, within which spending and frequency limits, using which data, and when must the assistant stop and ask?

Product implication: Treat broader authority as a new offer to evaluate. Do not infer it from previous use, an accepted recommendation or a general agreement to try AI.

Diagnose what “trust” means in the task

A single trust score can conceal several different concerns. For product research, separate at least four questions:

- **Capability:** Can the assistant complete this task accurately under realistic conditions?
- **Alignment:** Will its choices serve the consumer's stated preferences and interests?
- **Control:** Can the consumer constrain, correct and stop its actions?
- **Recovery:** What happens when the result is wrong or the consumer disputes it?

This is an organizing framework, not a validated psychological scale. Its purpose is to distinguish changes a business could make.

If people cannot verify a price or availability claim, improve the information and measure verification effort. If they suspect recommendations favor a sponsor, investigate commercial incentives and make relevant relationships intelligible. If they misunderstand the spending scope, redesign authorization and check comprehension. If they fear being unable to resolve an error, improve the actual recovery process and test it.

Provider familiarity may influence these judgments, but a recognizable brand is not evidence that an authorization design is understood. Similarly, a general promise of security should not be treated as proof that a particular data practice is acceptable to users.

Research should separate these factors where the business decision requires it. For example, hold task performance and price constant while varying one authorization design. If several product attributes change together, an observed difference cannot identify which one mattered.

A design that increases permission granting while worsening understanding is not a clear win. Evaluate whether users can correctly identify what the assistant may do, which limits apply and when further approval is required.

Product implication: Link each proposed change to a diagnosed obstacle and an observable result. Confidence-building language alone cannot establish that the product earned appropriate reliance.

Make control useful without making the work harder

Checkout.com's release also identifies spending caps, instant revocation and easy cancellation as prominent stated requirements. That makes them useful candidates for product evaluation; it does not establish their causal effect on permission or retention. [Checkout.com research](#).

Experimental evidence provides a narrower basis for taking control seriously. In three incentivized forecasting studies, Dietvorst, Simmons and Massey found that participants were more likely to select an imperfect algorithm when allowed to adjust its forecasts, even under restricted adjustment. The work concerns forecasting, not contemporary shopping agents. It supports testing meaningful control in the target task, rather than assuming the same effect transfers. [Published study](#).

For a shopping assistant, meaningful control may include specifying acceptable substitutions, approving changes above a threshold and withdrawing authority before another purchase. The benefit depends on whether the controls work and whether people can use them with reasonable effort.

Repeated confirmation can consume the convenience the product promises. Too little confirmation can create unwanted activity. The decision is where approval adds value relative to its burden, not whether every action should be automated or manually approved.

A suitable experiment could compare a small number of workable designs for the same routine task: approval for every order, standing permission with exceptions, and preparation without purchase authority. Keep demonstrated capability and relevant commercial terms comparable. Measure time and corrections as well as permission.

Revocation and cancellation must be distinguished. Revoking future authority need not cancel an order already placed. A cancellation request may also differ from a completed refund. State those boundaries and test whether users understand them. Describe only the recovery steps the service can actually deliver.

Product implication: Evaluate control as part of the service's usefulness. Avoid choosing a design solely because it removes the most clicks or obtains the most permissions.

Treat permission withdrawal as evidence

ACI's survey reports that 60% of UK adults said they would stop using an AI agent after one mistake. This is a stated reaction, not an observed churn rate. It motivates studying the experience after an error, rather than forecasting retention from that percentage. [ACI announcement](#).

Errors are not interchangeable. A late delivery, unwanted substitution, incorrect charge and breach of a spending limit can imply different product problems. Record the incident, its consequences, the consumer's understanding of it and the recovery outcome.

Withdrawal can indicate a sensible boundary. Someone may revoke standing permission but continue to use the assistant for comparisons. Treat that as a change in the relationship, rather than automatically labeling the consumer a lost user or an irrational skeptic.

For a product evaluation, observe repeated opportunities to use the service. Track initial permission, actual execution, consumer intervention, renewal or narrowing of scope, revocation and subsequent use. Retention should use a defined cohort and disclose whether users had another relevant task to complete.

If testing recovery, compare actual support arrangements under bounded, ethically reviewed conditions appropriate to the study. Do not infer the effectiveness of a refund promise from a survey that merely describes one. Distinguish hypothetical choices, incentivized task behavior and real transactions in the results.

Consumer protection and operational reliability are prerequisites, not attributes to remove merely to find a higher conversion rate. A design that preserves participation by obscuring cancellation or revocation would not support a defensible product recommendation.

Product implication: Measure whether the service remains useful after corrections and exceptions, and whether consumers can reduce its authority without unnecessary friction.

What evidence should change the product decision?

The premium research opportunity is to distinguish competing product actions. Public data has already identified concern about control and task differences. The missing evidence is what a specific design changes for a defined population under realistic conditions.

For routine replenishment, establish the intended decision before collecting data. Specify the product alternatives, target users, task conditions, acceptable quality, economics and conditions for expansion. The research must allow each design to succeed, fail or prove insufficiently useful.

Product decision	Evidence needed	Action the evidence can support
Select the initial task	Relevant demand, task performance and understood consequences	Choose a bounded use case worth testing
Choose the authority arrangement	Deliberate permission, comprehension and successful outcomes	Offer advice, approved execution or standing permission
Select the controls	Behavioral effects, usability and enforcement	Keep controls that improve appropriate reliance
Expand data access	Task benefit, understood scope and actual acceptance	Request only access justified by the service
Expand after initial use	Repeat value, incidents, recovery burden and economics	Expand, narrow, redesign or retain the existing scope

Measure both adoption and the cost of obtaining an acceptable result. Useful outcomes include preference fit, valid authorization, time spent reviewing, corrections, unwanted transactions and support effort. For the commercial case, include operating costs, cancellations, refunds and the contribution associated with completed activity.

Report assignment as well as actual use. Analyze relevant outcomes for the assigned population and disclose selective participation. A study restricted to enthusiastic early adopters answers a different question from one about a broader eligible market.

Do not use stated willingness as the sole go/no-go measure. Where feasible, test an implemented service with consequential choices. If access only permits a survey or controlled task, bound the recommendation to what that setting establishes and specify the field evidence still needed.

Set thresholds that fit the business decision and the user's interests. Higher conversion cannot compensate for unacceptable unauthorized activity. Greater retention may not justify support costs that undermine the service. If results are uncertain enough to reverse the decision, identify the next evidence needed.

The central recommendation is to build a **permission map** for the product: which task and authority arrangement each user accepts, under which limits, and how that scope changes with experience. This connects research to product choices more directly than one company-wide AI trust metric.

Method and limitations

This brief selectively reviews primary public sources checked on October 3, 2026. It is not an original participant study, systematic review, pooled estimate or reanalysis of respondent-level data. It does not identify a proven optimal authorization design or predict conversion for a particular product.

The evidence includes self-reported use, trust and hypothetical willingness, alongside older experiments on algorithmic forecasting. These measures do not establish real purchasing behavior. The U.S., UK and six-market findings concern different populations and instruments; they should not be averaged or compared as a league table.

Commercial survey releases summarize evidence without all the detail needed to assess recruitment, question order, weighting or statistical uncertainty. Gartner discloses its survey size and period; ACI discloses the YouGov sample, fieldwork and population weighting. Checkout.com's cited release identifies six consumer markets but does not supply a complete sampling and weighting account. No significance claims are made for the plotted differences.

ACI Worldwide and Checkout.com operate in payments, and Gartner sells research services. Their commercial interests provide context for their releases. This brief uses the reported findings without adopting promotional forecasts or causal claims.

The frameworks, recommendations and proposed evaluation are Eldris interpretations. Their application beyond shopping requires task-specific research. All figures distinguish reported survey measures from original conceptual guidance. No proprietary participant data or simulated observations are presented.

Sources

1. **Pew Research Center. Americans and AI 2026: Chatbots, Smart Devices and Views on Impact.** Published June 17, 2026; American Trends Panel survey of 5,119 U.S. adults, February 17-23, 2026. The cited measure is reported chatbot use. [Report, questionnaire and methodology.](#)
2. **Gartner. Gartner Survey Finds Consumers Want AI Shopping Help, But Not AI Purchase Decisions.** Published May 27, 2026. The cited willingness findings concern 322 U.S. consumers surveyed in January 2026; the release also covers a separate survey not used here. [Research release.](#)
3. **ACI Worldwide. Six in Ten UK Consumers Would Stop Using an AI Shopping Agent After One Mistake, ACI Survey Finds.** Published June 29, 2026; online YouGov sample of 2,080 UK adults, June 19-22, 2026, weighted to represent UK adults. [Issuer's announcement via Business Wire.](#)
4. **Checkout.com. Consumer demand for AI shopping is forming fast but trust for agentic commerce is still catching up.** Published June 9, 2026. Selected six-market consumer results; willingness and requirements are stated preferences. [Research release.](#)
5. **Dietvorst, B. J., Simmons, J. P., and Massey, C. Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them.** *Management Science*, 64(3), 2018, pp. 1155-1170. Three incentivized forecasting studies; not a study of shopping agents. [Published paper.](#)