

When AI Becomes the Default: Exposure, Choice and Dependence

An executive guide to deciding whether an AI feature should be automatic, opt-in or explicitly invoked

Consumer AI Adoption Gap · Research brief 07
October 5, 2026

© 2026 Eldris Inc. All rights reserved.

Executive brief

An AI feature can reach millions of people without anyone seeking out an AI product. A summary appears above search results. A suggested reply appears in an inbox. A familiar application gains an assistant. Distribution makes the feature available; product analytics must establish what users do with it and whether it helps.

Automatic exposure is evidence of reach. It is not, by itself, evidence of preference, trust or value.

Repeated exposure can come from repeated use of the surrounding product. A person may benefit from AI without clicking it, accept an output without understanding it, or continue using a product while avoiding its AI features.

This report is for consumer product leaders, research teams and AI providers choosing an activation design for one feature. Its bounded example is an AI summary inside an existing information-search product. The decision is whether the summary should appear automatically, require one-time opt-in or be requested for each task.

The central recommendation is to evaluate activation designs against **successful user outcomes and meaningful choice**, while measuring exposure separately. A default may reduce useful friction. It may also increase irrelevant output, conceal dissatisfaction or make weak demand look strong. Those possibilities require evidence from the feature and population being considered.

Three principles follow:

- Separate what the product displays from what the user requests, applies and returns to voluntarily.
- Interpret clicks, dismissals and session endings in relation to the task. The same event can indicate success or frustration.
- Compare activation designs using consistent populations and outcomes, then check whether any improvement persists beyond initial novelty.

Distribution changes the meaning of adoption


For a standalone assistant, opening the product can indicate an intention to use an AI service. For an embedded summary, opening the surrounding product does not establish that intention. Both are commercially relevant, but they answer different questions.

Pew Research Center's March 2025 browsing study found that 58% of respondents conducted at least one Google search associated with an AI summary. Separately, 13% visited an AI-tool website. The measures overlap and do not form a conversion funnel. [Pew browsing findings](#).

CONSUMER AI ADOPTION GAP / RESEARCH BRIEF 07

Exposure and deliberate destinations differ

Share of respondents · At least once in March 2025 · Pew Research Center



Different behaviors. Overlapping groups. Not a conversion funnel.

900 U.S. adults; tracked browsing. Neither measure establishes preference.

Source: Pew Research Center, May 23, 2025.

© Eldris Inc

Figure 1. Published respondent-level measures from a study of 900 U.S. adults. These observations do not establish that a person read a summary or preferred AI. Browser data exclude other apps and untracked devices. [Methodology](#).

The strategic implication is that distribution and demand can become difficult to separate. A feature inside a frequently used product can generate many impressions while attracting little deliberate attention. Conversely, a quiet feature can improve a task without becoming a destination users describe as an AI product.

For a product team, the question is not simply how many people encountered the feature. It is what the feature contributed to the existing experience. For an AI provider, embedding can open distribution opportunities, but the evidence needed to demonstrate customer value must follow the actual use case.

Decision implication: Keep reach, requests and outcomes in separate measures. Avoid comparing an automatically displayed feature with a deliberately opened assistant as though both numbers describe the same behavior.


A click is not a complete account of value

Pew's follow-up analysis reported clicks on traditional search results in 8% of visits associated with an AI summary and 15% of visits without one. This is an observed association, not a randomized estimate of the summary's effect. [Pew search analysis](#).

PUBLISHED OBSERVATIONS / NOT A RANDOMIZED COMPARISON

Traditional-result clicks differ across visits

Share of visits with a traditional-result click · Google searches



The next click does not establish task success or satisfaction.

March 2025 browsing; matching results retrieved in April 2025.

Source: Pew Research Center, July 22, 2025.

© Eldris Inc

Figure 2. Published visit-level percentages. Browsing was recorded in March 2025; matching search results were collected April 7-17, 2025. Results could change between dates. The chart does not establish task success, preference or causation.

Fewer clicks could mean that the summary answered the question. They could also mean that the person abandoned the task, accepted an incorrect answer or continued elsewhere. More clicks could reflect successful exploration or difficulty finding a useful answer. Telemetry records a behavior; interpreting the outcome requires more information.

For an informational task, combine behavior with an independently assessable outcome: did the person find the correct information, understand important qualifications and complete the intended task? For a shopping task, distinguish browsing, purchase and later regret or correction. Product retention alone cannot settle these questions.

Different businesses also face different decisions. A search provider may care about useful answers and continued use; a publisher may care about qualified referral traffic. A retailer may care about appropriate purchases and returns. A metric can improve for one participant while worsening for another. State whose outcome the decision is intended to improve.

Decision implication: Treat engagement measures as evidence to explain. Choose task outcomes that can distinguish efficient completion from abandonment or misplaced confidence.

Measure five different relationships with the feature

The framework below is original Eldris guidance. It separates observations that often get compressed into an adoption rate. It is not a required sequence: useful assistance can occur without explicit invocation, and continued exposure can occur without acceptance.



Figure 3. A conceptual measurement map. Each observation supports a different claim. None alone establishes overall product value, and the rows are not stages in a conversion funnel.

Exposure

Record when a feature is rendered and when it is actually visible. A generated summary below the visible area is different from one displayed prominently. Neither establishes attention or comprehension. Keep people, sessions and task occasions separate.

Deliberate invocation

Record an explicit request for the feature. This can indicate interest at that moment, but invocation also depends on discoverability, effort and the task. Low invocation could reflect limited demand or an interface that prevents people from finding useful assistance.

Application of the output

Identify whether the output enters the user's work or decision. For a writing feature, insertion and subsequent edits can be observed. For a summary, reading and using the information may require task-based observation. Leaving an output unchanged does not establish accuracy or understanding.

Continued use when choice is available

Examine repeat use on later relevant tasks when an alternative remains accessible. Persistence of a default setting is weak evidence of active preference: people may not know how to change it or may not consider the change worth the effort. Interpret reuse alongside dismissals, disabling and qualitative explanations.

Reliance

Ask whether users now depend on the feature to complete the task, and on what basis. Test understanding of limitations and response to conflicting information. Reliance can be appropriate when the feature delivers dependable value. It can also reflect habit, loss of alternatives or confidence unsupported by performance.

For the proposed summary feature, success might include accurate answers, lower effort, appropriate source checking and later voluntary use. The team should choose a small set of measures that fit the task rather than instrumenting every possible interaction.

Decision implication: Make the claim attached to each metric explicit. A dashboard should not relabel exposure as use or repeated use as informed reliance.

Choose an activation design for the task

Automatic, opt-in and explicitly invoked designs distribute effort and choice differently. None is universally superior.

Automatic display makes assistance available without an extra request. It may help users who would not discover the feature independently. It also creates outputs for people who do not want them or for tasks where they add little. Assess distraction, latency, errors and the ease of ignoring or disabling the feature alongside any benefit.

One-time opt-in asks people to enable the feature before routine display. It provides an observable initial choice, but that choice can become stale as the feature or its permissions change. Evaluate whether users understand what they enabled and can revise the choice.

Explicit invocation for each task allows selection at the point of need. It adds interaction effort and depends on discoverability. It may suit situations where users can judge when help is useful or where unwanted outputs are disruptive.

A reasonable product may combine these approaches: automatic display for a narrow category and invocation elsewhere. The boundaries should follow evidence about task suitability rather than a desire to place AI everywhere.

Human-AI interaction research provides useful design guidance on communicating capability, facilitating invocation and dismissal, supporting correction and giving users control. These principles help frame the evaluation; they do not establish which activation design wins for a particular feature. [Human-AI interaction guidelines](#).

Decision implication: Compare the complete user experience of the alternatives, including the additional steps, unwanted output and ability to change course. Hold core capability constant where the purpose is to isolate activation design.

Run a comparison that answers the rollout decision

For the search-summary example, define eligible tasks and users before comparing designs. Include the existing experience without the summary as a baseline when feasible. A comparison among three AI variants alone may miss the possibility that none improves the task.

A product experiment can assign eligible users to activation designs and observe subsequent task outcomes. Maintain a consistent experience for each assigned user when learning or settings persistence would otherwise blur the comparison. Define the assignment unit, observation period, primary outcome and acceptable limits before examining results.

Evaluate the population assigned to each design, including people who never invoke the feature or choose not to enable it. Comparing only active AI users would select different people under each design and could conceal the real rollout effect. Analyses of affected task occasions can add detail, but eligibility must be defined consistently rather than by treatment-induced engagement.

Microsoft's experimentation guidance explains how changes in denominators and poorly defined affected-user segments can distort interpretation. This is a methodological foundation, not evidence for an AI-specific product recommendation. [Experimentation guidance](#).

Report both outcomes and exposure. One design may improve completion for a small group while degrading the experience of many others. Examine task categories or user groups identified before analysis where there is a credible reason to expect different effects. Do not invent a winning segment after an inconclusive result.

Observe more than initial curiosity. Repeat measurement on later relevant tasks, accounting for people who leave or disable the feature. Choose an observation period that fits how frequently the task occurs. A single paid research session can reveal usability or immediate task performance; it cannot establish durable reliance in everyday use.

Decision implication: Select the design based on the intended rollout population and meaningful outcomes. Use usage among self-selected adopters to understand behavior, not as a substitute for the rollout comparison.

Distinguish valued reliance from constrained dependence

A feature can become difficult to avoid because the product makes alternatives inconvenient, because an organization mandates the surrounding software, or because the user loses practice with an unassisted task. Continued activity under those conditions needs different interpretation from voluntary preference.

Conversely, automatic assistance can be genuinely valuable even when users do not want to manage another setting. The aim is not to require constant deliberation. It is to establish that the feature serves the task and that relevant choices remain understandable and usable.

For the summary, observe what happens when source material contradicts it. Can users recognize the discrepancy and complete the task? Do they check when the consequence warrants it? Can they dismiss the summary or use the original results without losing their place?

A planned comparison with the unassisted experience can reveal whether the feature adds value and whether important skills or comprehension are being bypassed. Such comparisons should preserve an appropriate way to complete the task; disruption caused by removing a familiar interface is not automatically evidence that AI is indispensable.

Decision implication: Pair continued-use measures with outcome quality, understanding and the practicality of alternatives. Commercial retention is useful evidence, but it does not establish that reliance is well founded.

What to require before making AI the default

Prepare a short decision record:

1. **Task and population:** who receives the feature, on which occasions and for what purpose.
2. **Activation alternatives:** automatic, opt-in, invocation and the existing experience, with differences in capability or friction stated.
3. **User outcome:** how successful completion, effort and important errors are measured.
4. **Choice and control:** whether users can understand, dismiss, correct or disable the feature without unnecessary disruption.
5. **Comparison evidence:** results for consistently defined populations, uncertainty, observation period and relevant subgroup limits.
6. **Decision and review:** the chosen scope, reasons, unresolved conditions and triggers for reconsideration.

Automatic display is supportable when evidence shows worthwhile outcomes within the chosen limits. Opt-in or invocation may be preferable when benefit is selective, tasks vary or unwanted assistance creates material friction. Preserve the existing experience when the evidence does not justify the change.

For AI providers, the commercial opportunity is to demonstrate value inside the partner's actual workflow. Impressions can show distribution. Evidence of better tasks and valued repeat use makes a stronger product case.

The practical test is straightforward: **Would we still choose this activation design if the dashboard could not count an automatically displayed output as an AI user?** If the answer changes, the current argument may depend more on the measurement than on the experience.

Method and limitations

This brief selectively reviews primary sources checked on October 5, 2026. It is not an original experiment, systematic review or current estimate of consumer AI adoption. Pew's observations concern March 2025 browser activity; its two analyses use the same panel and are not independent replications. They do not establish causation, satisfaction or informed reliance.

The five-observation map, summary-feature example and decision record are original Eldris analysis. No product results or participant data are claimed. Interaction and experimentation sources inform design and measurement; they do not demonstrate the effects of the activation alternatives proposed here.

Public evidence supports the distinctions in this report. Choosing a default for a specific product requires feature-specific evidence. Task studies can clarify immediate behavior; field comparisons and follow-up are needed to assess routine use and continuing reliance.

Sources

1. **Pew Research Center. What Web Browsing Data Tells Us About How AI Appears Online.** May 23, 2025. March 2025 browsing study; respondent-level exposure and AI-tool website visits. [Report](#).
2. **Pew Research Center. Google users are less likely to click on links when an AI summary appears in the results.** July 22, 2025. Observational follow-up using the same panel; traditional-result clicks by summary presence. [Analysis](#).
3. **Pew Research Center. Methodology: What Web Browsing Data Tells Us About How AI Appears Online.** May 23, 2025. Panel selection, consent, tracking and weighting. [Methodology](#).
4. **Amershi and colleagues. Guidelines for Human-AI Interaction.** CHI 2019. Design guidance evaluated across interaction scenarios; it does not compare the activation designs in this brief. [Publication and paper](#).
5. **Microsoft Research. Patterns of Trustworthy Experimentation: Post-Experiment Stage.** Methodological guidance on denominator changes, affected-user analysis and interpretation of rollout effects. [Guidance](#).