

Fall 2026

F-Secure Digital Trust Report

What Consumer Behavior Around AI
Reveals About Digital Trust



Using an AI Chatbot Isn't the Same as Trusting It

“AI chatbots occupy a strange place in everyday life. They are increasingly familiar, frequently useful, and widely questioned. People turn to them for quick answers, advice, and decisions, yet remain uneasy about whether their responses can be trusted.

Using AI chatbots doesn't mean that users trust them, especially in an age where personal information is harvested everywhere. Every app tracks something, every site asks for consent, and every scroll gets logged. AI chatbots sit at the center of that system and ask for more than anything else: not just what you click, but what you think.

The danger is that users are sharing private parts of their lives with AI chatbots. We found that over a third (33.9%) of users are asking AI for personal advice on health, money, and relationships — the exact categories that people are most guarded about everywhere else.

That's why simply asking, “Do you trust AI?” isn't enough.

Trust is so complex. It's situational and contextual, meaning it changes with the task, the stakes, and the moment. Someone may distrust AI in principle but accept its answer when it sounds convincing, when the question feels simple, or when checking would take too much time.

What people say about trust and what they do with an answer can be different.

We surveyed 1,500 consumers in the US and UK to examine this gap. Instead of treating trust as a single attitude, we looked at the behaviors that surround it: when users verify an answer, what prevents them from checking, what makes them accept a response, and whether habits change as they use AI in more areas of life. We also asked non-users what keeps them away and whether stronger security would change their minds.”



Dr Amel Bourdoucen

Senior User Researcher,
Human-Centric Cyber Security and Privacy
F-Secure

What's in the Report

01 – The More Ways People Use AI, the Less They Verify Answers

02 – For Almost Half of AI Users, “Sounding Right” is Good Enough

03 – Security Should Address the Risks Consumers Care About

04 – Trust in the AI Era: Key Recommendations for Digital Service Providers

KEY FINDINGS

- **9 in 10** AI users say checking answers is their responsibility, yet 70.1% of them don't check consistently.
- **41.4%** of broad AI users rarely or never check answers, compared with 25.8% of those who limit AI to certain tasks.
- **47.3%** of AI users don't question an AI chatbot's answer further when it sounds right.
- **76.6%** of AI users have engaged in at least one risky AI behavior.
- **46.1%** of non-users say scam detection would matter most in AI chatbot security — the most selected feature.


01 – The More Ways People Use AI, the Less They Verify Answers

Using AI involves a degree of uncertainty. Chatbots can provide useful and convincing answers, but they can also hallucinate, generating fabricated information and presenting it as fact. So, when do people decide a response needs checking, and when do they simply take it at face value?

Who people say is responsible for checking AI answers



How often people actually check AI answers



Roughly 9 in 10 (89.5%) AI users say they are fully or mostly responsible for checking AI responses. Yet 70.1% of this group don't check consistently, doing so only sometimes, rarely, or never.

That contrast is even more striking given that almost half (49%) believe it's fully their responsibility to check AI answers. People accept responsibility for verification in principle, but often don't follow through in practice.


SURVEY RESPONDENT

"It generally gives quick and accurate responses, but it's my responsibility to fact-check everything and look for biases."

Male, 35-44, US


Base: AI users who say checking is fully or mostly their responsibility, n=933 (responsibility base n=1,045). Source: US and UK AI survey, April 2026 (Q81, Q71).

What stops people from checking AI answers



Base: AI users, n=1,045. Multi-select, sums above 100%. Source: US and UK AI survey, April 2026 (Q76).

Where people check AI answers




Base: AI users, n=676. Multi-select, sums above 100%. Source: US and UK AI survey, April 2026 (Q72).

When people skip checking an AI answer, the biggest barrier isn't effort. It's having no reliable source to check against (37.5%), followed closely by the answer already seeming good enough (34.4%). In other words, people are less likely to check when they lack a clear reference point, or when the answer itself gives them little reason to doubt it.

Time and ability aren't what hold most people back.

When people do check an answer, most look beyond the original chatbot: more than three-quarters (76.8%) search online and over half (53.4%) draw on their own knowledge. But some turn back to AI. A quarter (25.4%) ask another chatbot and one in five (20.1%) ask the same one, effectively using AI to check AI rather than independently verifying the original answer.

How checking changes with broader AI use



Base: AI users who answered the checking question, n=1,043. Association only, not causation.
Source: US and UK AI survey, April 2026 (Q68, Q71)

How broadly people use AI is more strongly linked to checking behavior than how often they use it. Among broad users (people who reach for AI across a wide range of tasks) 41.4% rarely or never check clear, confident answers, while only 24.7% check regularly. Among scoped users (people who limit AI use to certain tasks) the pattern reverses: 25.8% rarely or never check, while 36.8% check regularly.

As AI becomes a tool for more areas of people’s lives, they appear to question its answers less consistently.

One possible explanation is that as AI use broadens, task-by-task judgment may give way to a more generalized sense that the tool is reliable.

But the relationship could work in either direction: broader AI use and less frequent checking are linked, but we can’t say which comes first.

What the data tells us

People don’t appear to treat verification as a consistent step in using AI. Instead, they are more likely to skip it when an answer already seems convincing or when they have no reliable source to check against.

Broader AI use is also linked to less consistent verification, despite most people believing they are responsible for checking AI answers. The challenge isn’t knowing that AI can be wrong. It’s knowing when an answer needs checking and having something trustworthy to check against.

What this means for digital service providers

Providers shouldn’t assume consumers will consistently recognize when an AI answer needs checking. As AI becomes embedded in more digital experiences, such as customer service chatbots, AI assistants, and website search, creating trustworthy experiences means helping consumers navigate uncertainty rather than simply asking them to be more careful. This could look like clearer, more reliable sourcing and confidence indicators that show when information can or cannot be verified, potentially helping consumers make informed decisions.

02 – For Almost Half of AI Users, “Sounding Right” is Good Enough


If people know AI can be wrong, why are they not questioning its answers?

You might expect it to be time pressure, fatigue, or a lack of effort. However, our data points to the answer itself. Almost half (47.3%) of AI users say they don’t question an answer further when it sounds right, more than double any other behavior surveyed.

More than 9 in 10 (90.5%) AI users recognize at least one of these behaviors in themselves, but “sounding right” stands apart. While the other responses relate to the user’s time, effort, or attention, it’s the answer itself that gives them a reason to stop questioning it.

And for most people who selected it, that reason appears sufficient on its own: 61.2% reported no other behaviors alongside it. This suggests an AI answer doesn’t always need to be proven right; it only has to sound right enough to end the search.

What people recognize in their own AI use



SURVEY RESPONDENT

“It’s like having a super-fast research assistant who’s amazing at brainstorming but occasionally gets a bit too confident about facts they’ve made up.”

Female, 25-34, UK


Base: AI users, n=1,165. Multi-select, sums above 100%. Source: US and UK AI survey, April 2026 (Q74).

Who is most likely to take a convincing answer at face value? Our data points less to how often someone uses AI and more to how widely they use it.

Acceptance of answers by how often people use AI



Acceptance of answers by how widely people use AI



Base: AI users answering both questions, n=1,165. Selected “If it sounds right, I don’t question it further” (Q74). Association only, not causation. Source: US and UK AI survey, April 2026 (Q74 with Q40, Q68).


How often people use AI matters, but only up to a point. Acceptance of responses rises sharply from 44.1% for weekly users to 56.5% for daily users, then barely changes for those using AI several times a day (58.1%).

How widely people use AI shows a much bigger difference: 28.6% of people who keep AI to low-stakes tasks take answers at face value, compared with 59.7% of those who use it for many kinds of tasks.

People who use AI across more areas are far more likely to accept its answers.

Accepting an answer because it sounds right is also part of a broader pattern of trust and reliance on AI. This pattern goes beyond accepting answers. People are using AI for verification, turning to it for help with important decisions, and sharing potentially sensitive information with it.

How people rely on AI



Base: AI users, n=1,165. Multi-select, sums above 100%. Source: US and UK AI survey, April 2026 (Q73).

These behaviors often overlap. Over three-quarters (76.6%) of AI users report at least one, and among those people, more than 7 in 10 (73%) report two or more, showing how trust and reliance on AI can take multiple forms.

What the data tells us

People don't always stop questioning AI because checking takes too much time or effort. Often, the answer itself is

enough: once it sounds plausible enough, further scrutiny can feel unnecessary.

This tendency is more common among people who use AI across many kinds of tasks. The risk is that sounding right can become a substitute for establishing whether an answer is actually right.

What this means for digital service providers

For providers using or considering using AI chatbots as part of their service, this means reducing the 'marginal cost' of checking each answer. Consumers shouldn't need to conduct a separate search every time, particularly when the right source for verification may depend on the context and their own understanding of the topic.


Further research work in this area must explore what sits behind the 76.8% of users who searched online and the 25.4% who asked another chatbot. This means looking more closely at users' mental models and perception of using these tools for verification, because it's not always clear where 'traditional search' ends and 'AI search' starts.

Tools like Google AI Overviews surface AI-generated summaries directly within search results, so a user who is 'searching online' may in fact be reading an AI-generated answer without recognizing it as such. Our upcoming consumer research will explore this further.

03 – Security Should Address the Risks Consumers Care About

What role can stronger security play for people who don't currently use AI chatbots? The answer depends partly on why they don't use them in the first place.

Why people don't use AI chatbots



Distrust (42.9%) and not seeing the need (42.5%) are virtually tied as the two most common reasons people don't use AI chatbots. Despite the similarity in size, they represent fundamentally different barriers.


Someone who doesn't trust AI has a reason not to use it. Someone who doesn't see the need has no reason to start.

This matters when considering what stronger security can and can't address. Security can respond to concerns about risk, but it can't resolve every reason people have for not using AI.

So, we asked whether a built-in security layer that warned about misleading content and risky data sharing would make non-users more likely to use an AI chatbot.


Base: Non-users of AI chatbots, n=280. Multi-select, up to three choices, sums above 100%. Source: US and UK AI survey, April 2026 (Q85).

Would stronger security make non-users more likely to use AI?



Base: Non-users of AI chatbots in the last 3 months, n=280. Single select answer. Source: US and UK AI survey, April 2026 (Q88).

How different non-users respond to stronger security



Base: Non-users citing distrust only, n=73; non-users citing no need only, n=72. Respondents citing both reasons (n=47) are excluded. Difference in “still wouldn't use”: p<0.05. Source: US and UK AI survey, April 2026 (Q88 x Q85).

Stronger security does increase willingness to use AI for a sizable minority of non-users. More than a third (36.5%) say a security layer would make them at least slightly more likely to use an AI chatbot. At the same time, more than half (51.4%) say they still wouldn't use one, suggesting security alone can't overcome every barrier to adoption.

The response also differs depending on why someone doesn't use AI in the first place.

Among people whose only barrier is distrust, 24.7% say stronger security would make them more likely to use AI, compared with 33.3% of those who simply don't see the need. Security therefore has some influence across both groups, but it isn't enough on its own to overcome established distrust or create a compelling reason to use AI.

Willingness to adopt AI is only part of the picture. The features non-users prioritize tell us something different: which risks they most want security to address.

What non-users want from stronger AI security




Base: Non-users of AI chatbots, n=280. Multi-select, up to three choices, sums above 100%. Source: US and UK AI survey, April 2026 (Q89).

Only 16.1% cite scams as a reason they don't use chatbots, yet scam detection is the most requested security feature. This shows a disconnect, as the risks stopping someone from adopting AI aren't necessarily the same risks they expect security to protect them from.

More broadly, non-users favor active protection and warnings over information that leaves them to assess risks themselves. Scam detection (46.1%), preventing data from being stored or shared (45%), blocking harmful or manipulative content (42.9%), and flagging potentially wrong or misleading answers (37.5%) all rank well above providing sources or explanations for AI answers (14.6%) or a confidence or reliability score for responses (11.4%).

What the data tells us

Consumers have clear expectations of what AI security should do. Non-users prioritize active safeguards against scams, harmful content, risky data sharing, and misleading answers over explanations or confidence scores.



Stronger security won't persuade everyone to use AI, nor should that be its only measure of success. Security's value lies in addressing the risks consumers care about and providing protection as AI becomes embedded in more digital experiences, including those where consumers may not actively choose to engage with it.

What this means for digital service providers

Providers have an opportunity to make security part of the customer experience. The security feature preferences point toward the kind of protection that could make those experiences feel more trustworthy: safeguards that actively respond to scams, risky data sharing, harmful content, and misleading answers rather than leaving customers to interpret risk information themselves.


Security for AI chatbots is still an emerging research topic, and this is a broad, industry-level problem today. At F-Secure, we are actively researching this area, focusing on our understanding of these tools and how people perceive and use them. We intend to remain active in this space, working toward the protection consumers will need as these risks become clearer.

Trustworthy Systems Are Built, Not Claimed

“We keep asking whether people trust AI. But trust is what consumers feel; trustworthiness is what we build — and it’s the part a product builder or service provider can actually do something about. It comes down to three things: being honest about what a system can and can’t do, competent enough to do the job well, and reliable enough to do it again tomorrow. Those are design and engineering choices, not marketing ones. Our job isn’t to persuade people to trust AI. It’s to make things that are honest, competent, and reliable — and let trust follow.”

Dr Laura James

Vice President, Research
F-Secure



04 – Trust in the AI Era: Key Recommendations for Digital Service Providers

1. Leverage existing customer trust

The biggest barrier to checking AI answers isn't time or effort. It's having no reliable source to check against (37.5%). Providers can build on existing customer relationships by making AI experiences easier to verify through clear and reliable sourcing. That creates an opportunity to turn existing customer trust into a differentiator as AI becomes more embedded in digital services.

2. Design experiences around user behaviors

How widely people use AI is more strongly linked to verification than how often they use it. Among people who use AI across many kinds of tasks, 41.4% rarely or never check confident-sounding answers. As people rely on AI across more areas of their lives, providers can't assume they'll always assess the risks themselves. Protection should be built into the experience rather than relying on consumers to seek it out.

3. Prioritize intervention over information

Non-users rank scam detection (46.1%), preventing data from being stored or shared (45%), and blocking harmful content (42.9%) well above sources (14.6%) and confidence scores (11.4%). The message to providers is clear: consumer protection is necessary to navigate the emerging risks with AI chatbots today.





“We stand at a trust inflection point: one that will divide the companies who saw it coming and acted from those who didn’t. Which side of the line will you choose to be on?”

Timo Laaksonen
President and CEO
F-Secure

Methodology

We surveyed 1,500 consumers aged 18–64 across the US and UK in April 2026, including frequent and occasional AI chatbot users, as well as people who don’t use AI chatbots.



How to Read the Data

Bases

AI users: n=1,045–1,165 depending on the question asked.

Non-users: n=280, defined as everyone who directly answered “No” to Q38 (“Have you used any AI tools or chatbots in the last 3 months?”). Exact bases are shown below each chart.

Multi-select questions

Some questions allowed more than one answer, so percentages may add up to more than 100%.

Comparisons

Relationships between groups show association, not cause and effect.

Statistical confidence

Where statistical significance is reported, comparisons use a 95% confidence level.

Geography

US and UK responses are combined throughout unless otherwise stated.



“

“Illuminate, F-Secure’s research function, brings together experts to explore the human, social, and technical aspects of security. We identify emerging threats, prototype new protection systems, and anticipate future risks to keep consumers safe. By staying ahead of the curve, we navigate a constantly evolving digital world and ensure F-Secure delivers trusted, reliable, and innovative cyber security solutions.”

Dr Laura James

Vice President, Research
F-Secure

About F-Secure

F-Secure is a human-first, AI-powered consumer cyber security experience company with 38 years of expertise in tackling digital threats. We help Digital Service Providers turn trust into a high-value growth engine — protecting their customers while enabling them to live their best digital lives in a world of relentless, AI-driven scams.

With billions of digital interactions secured each year, tens of millions of consumers protected globally, and over \$10bn in partner value created, we deliver proven impact at scale.

To find out more visit f-secure.com/partners or follow us on our social channels.

