

The State of Consumer AI Research Report 2026

Understanding AI Usage and What That Tells Us About the Future of
Personal Intelligence

A study of 1,000 chatbot users and 100 non-users

Jofish Kaye, Ph.D.

Inflection AI

July 2026

Contents

Executive summary	3
1. Introduction	4
2. Methodology	4
Chapter 1 — The everyday chatbot user	6
3. The everyday chatbot user	7
4. What people use chatbots for	8
5. How people interact	9
6. What people want chatbots to become	11
7. How chatbots are rated	12
8. Paying for chatbots	13
9. Why people walk away	14
10. If it were a dinner guest	15
Chapter 2 — Five kinds of chatbot users	17
11. How people choose a chatbot	18
12. Five kinds of chatbot users	19
Chapter 3 — The frustration is real	22
13. The people who don't use chatbots	23
Chapter 4 — Conclusions	29

Executive summary

This report draws on two surveys conducted in March 2026: one of 1,000 U.S. adults who had used an AI chatbot in the previous week, and a companion study of 100 adults who had not. Together they describe who uses these tools, how and why they choose between them, and — just as importantly — why a substantial minority continue to keep their distance.

Five findings stand out:

- **Chatbot use is a healthy, plural ecosystem.** Users report using an average of 4.2 chatbots and reaching for roughly 1.8 of them daily. The market looks less like the two-horse race of mobile operating systems and more like podcasting: a few large players alongside a vibrant long tail.
- **People want chatbots to do more than think for them.** Asked which future roles would be most valuable, respondents favored current productive roles — assistant, teacher, editor — but 80% also chose at least one relational or supportive role (mentor, coach, listener, companion and the like), and 55% chose more than one.
- **Even in jest, people perceive chatbots as having personalities.** Asked to picture each tool as a dinner guest or an animal, respondents produced vivid, consistent archetypes — the “popular professor,” the “edgy rebel,” the loyal golden retriever — a sign of the emotional relationship many have with these tools.
- **There is no single “chatbot user.”** Cluster analysis yields five segments, from outcome-focused Capability Seekers to relationship-driven Emotionally-Led Power Users, each with a different demographic profile and a different vision of what a chatbot should become.
- **Non-use is driven by distrust as much as by absence of need.** Asked in their own words, most non-users simply have no use for a chatbot. Handed a checklist, they reach overwhelmingly for ethical and environmental objections. The gap between the two says a great deal about how non-users think about AI, and the heart of Chapter 3.

The strategic implication runs throughout: the more sharply an AI assistant maker can identify who it serves, and the more honestly it listens to those it does not, the more useful a product it can build. No one needs to solve every problem for everyone.

1. Introduction

Almost everyone has an opinion about who uses AI assistant chatbots. Far fewer people have asked. So we did.

For anyone building these tools, today's users are valuable but no longer the whole story, and meaningful growth means looking further afield. One large and reachable opportunity is people who already use chatbots: persuading them to add another tool to a list they already maintain is a smaller ask than changing the habits of someone who uses none. That population is the focus of the primary study in this report. A second, smaller study turns to the harder and equally instructive question of people who don't use chatbots at all. Identifying the characteristics that make a chatbot attractive to a particular group is what lets its makers focus effort where it will matter most.

This report is organized in four chapters, mirroring the arc of the research. Chapter 1 describes the everyday chatbot user — who they are, what they use these tools for, how they interact, what they want them to become, what they pay, why they sometimes walk away, and the personalities they project onto them. Chapter 2 turns to how people choose between chatbots, and the five distinct segments those choices reveal. Chapter 3 examines non-users, and argues that their reasons — and the broad distrust that lies beneath them — deserve to be taken seriously rather than treated as obstacles to be overcome. A short fourth chapter draws the conclusions together and sets out where the research should go next.

2. Methodology

The chatbot-users study (N = 1,000)

After an initial pilot survey of 100 chatbot users which we discarded, we recruited 1,000 respondents through Prolific, all primary English speakers living in the United States, in a sample matched to the U.S. census on sex (51% female, 49% male), age band, and geographic region. The median respondent was 46 years old and spent about 15 minutes on the survey, for an incentive of \$6.75. A sample of this size means roughly a 95% chance that our figures sit within about 3 points of the true value.

One caveat up front: a Prolific panel is, by definition, made up of people comfortable enough with the internet to take online surveys for money. The median age of the U.S. population is 39; our English-speaking, online sample skews a little older (46) and a lot more tech-forward than the country as a whole. That said, our sample also likely excludes professional users, like developers with heavy Claude Code habits, who are unlikely to be answering survey questions for \$27/hour.

We screened carefully for this study. We didn't start by asking for people who used chatbots: respondents just knew that we wanted to ask them about their technology use. The first question asked them "Which of the following have you used in the past week?" with a list of some twenty possibilities, and we only continued with people who told us they had used "A chatbot, like

Claude or ChatGPT”. But, in passing, we learned some interesting things: 61% had listened to a podcast by themselves; 72% told us they had used wireless headphones this week; 16% told us they had listened to a podcast with other people.

Of about 1,300 people who began answering the study, we filtered out the 15% who hadn’t used a chatbot in the past week. We also planted honeypots — a fake “barcode-based household multimedia system” that 64 people claimed to use, and, later, a fictional chatbot called “Afitoto” that 3 people selected. Everyone who selected one of those was removed.

The non-users study (N = 100)

A companion survey, completed on 30 March 2026, recruited 100 U.S. adults through Prolific who had not used a chatbot in the previous week. The median respondent was 46 and spent about 12 minutes for an incentive of \$5.25. We recruited 743 respondents to find our 100 non users, in addition to filtering out another 43 respondents who told us they had used that “barcode-based household multimedia system” this week. With a sample of 100, the margin of error is close to 10 points, so its figures are best read as broad guidance rather than precise measurement. The same Prolific sampling caveat applies.

CHAPTER ONE

The everyday chatbot user

Who uses chatbots in 2026: their lives and devices, how they interact, what they want these tools to become, and the personalities they give them.

3. The everyday chatbot user

Compared to the general population, chatbot users are early and heavy adopters of consumer technology across the board. The gap between this survey and a general-population baseline is striking on almost every device category. A note on the comparison: the general-population figures throughout this report are not measured values but approximations generated by Claude for context, and should be read as rough benchmarks rather than precise statistics.




Figure 1. Weekly technology use: chatbot users versus the general population (general-population figures are Claude-generated approximations for comparison).

The takeaway isn't that chatbot users are exotic — it's that they live comfortably inside a connected ecosystem. They stream, they listen, they wear the watch. A chatbot is one more tool in a well-stocked drawer.

Most people use multiple chatbots

Most people aren't loyal to a single assistant. Respondents reported using or having regularly used an average of 4.2 different chatbots (median 4.0), and using about 1.8 chatbots “several times a day” or “at least once a day.” Heavy users pull that average up; the median daily count is just 1. And remember, this survey undercounts professional usage — the developer firing up an AI coding assistant all day for work isn't answering a survey for \$6.75 in return for 15 minutes of their time.

We think this is one of the most important results from this study. This is a healthy ecosystem, where many people are using multiple chatbots for different reasons. As an analogy, it looks more like podcasts — where there are a few big players, but there's a vibrant ecosystem which caters to different people and their interests — than it looks like mobile phones, where there's two big choices to pick between.


ChatGPT is the most popular, with 91% of respondents telling us they use or have used it regularly, followed by Gemini at 64% and Siri at 45%. In this crowd, people who currently use or have previously used Claude show up at 29%.

4. What people actually use chatbots for

It is easy to form an impression of chatbot use from what circulates in public: the screenshots people post, the threads they brag in, the questions they ask out loud on social media. That is a heavily filtered view, skewed toward the impressive and the shareable. It is tempting to draw conclusions from the posts people make about their AI usage, but the picture here is different. These are ordinary people telling us privately, in a survey, what they last did with a chatbot — not performing their use for an audience. And what they describe is overwhelmingly mundane, practical, and woven into everyday life.

We asked simply: “What was the last thing you used a chatbot for?” An approximate coding of the open-ended answers is shown below.

Figure 2. What respondents said they last used a chatbot for (approx



Approximate coding of open-ended answers, $n = 1,001$; responses may touch more than one theme.

The single largest category, by a wide margin, is quick information and search — roughly a third of answers — followed by work tasks, learning and explanation, and writing and editing. Cooking, health questions, personal finance, and coding each take a smaller but steady share. The individual answers are telling in their ordinariness: “how to figure out something on GitHub,” “a question to do with electricity,” “Bible study,” “help me form an investment strategy,” “a plan to alleviate some financial hiccups,” “important information about growing vegetables,” “information about non-alcoholic fatty liver disease,” “comments and descriptions

for my real estate listings.” This is the chatbot as a general-purpose helper threaded quietly through daily life — not a stage for impressive demonstrations.

What they wish a chatbot could do

We also asked whether there was something they wished a chatbot could help with but hadn’t found done well. The most common answer was a version of “not really” — many are satisfied with what they have, or are still new and exploring. Among those who named something, the wishes cluster less around dramatic new powers than around the existing thing simply working better: reliability and accuracy above all (“work on a long-term project without hallucinations or sputtering out”; “when I have a long conversation it tends to forget what it had previously [said]”), reaching real-world content like web pages and images, or occasionally reaching for something more human (“companionship”; “getting an actual person when that is what you need”). The gap people feel is mostly one of trust and dependability, not missing features — a theme that returns, far more sharply, among non-users in Chapter 3.

5. How people interact: typing, talking, phone and computer

We asked respondents, for each chatbot they used, exactly how they reached it: typing or talking, and on a phone, a computer, in the car, or through a smart speaker. Two dimensions emerge from the answers — whether people type or talk, and whether they use a phone or a computer — and each tells a different part of the story.

Start with input. For all the attention paid to voice assistants and ambient AI, the dominant mode of interaction remains the keyboard. Typing dominates across nearly every tool. The clear exceptions are Siri and Alexa, the only assistants with substantial voice and in-car use, where upwards of 95% of use is spoken and barely a fifth involves any typing at all.




Figure 3. Mode of interaction by chatbot, as a share of each tool's users.

This points to a real and largely unmet opportunity. No single chatbot does voice really well except Siri and Alexa — and at the time of running the study, neither was backed by a very powerful AI model. The most capable, highest-rated tools are almost never spoken to aloud, while the two assistants people do talk to are the least able to hold up their end of the conversation. That leaves a conspicuous gap between capability and modality: a genuinely powerful chatbot that is also pleasant to talk to is, for now, a space nobody has filled. Confidence is not the barrier — asked how comfortable they feel using new technology, respondents returned a median of 6 on a 7-point scale. This is a self-assured crowd that has simply not been given a compelling reason to talk to its most capable tools.

The device split is just as revealing, and it separates the chatbots by purpose. Because people could choose more than one setting, the shares below do not sum to 100 — many use the same tool on both a phone and a computer.




Figure 4. Phone versus computer use by chatbot, as a share of each tool's users.

The work-oriented assistants lean heavily to the computer: roughly 89% of CoPilot and Claude users reach them on a desktop, against about half on a phone. The opposite holds for the more conversational and mobile-native tools — Meta AI (85% phone), Grok (80%), and the phone-bound Siri (93%) — which live in people's pockets. The two market leaders, ChatGPT and Gemini, are unusual in being used heavily on both: around three-quarters of their users reach them on a phone and roughly four-fifths on a computer. The voice assistants are a category apart — Alexa's use is concentrated in the smart speaker rather than either a phone or a computer.

6. What people want chatbots to become

One of the most telling questions was this: "In the future, chatbots will be able to do lots of things better than now. Pick the five from this list that you think will be most valuable to you as chatbots improve and evolve."




Figure 5. Roles respondents chose among the five they would find most valuable (n = 1,001).

When we looked at the whole thousand-user population of this study, we saw some clear patterns. The most popular choices were the productive roles that chatbots play well now: an assistant, a teacher, and an editor, each picked by 65%+ of respondents. Surprisingly, perhaps, the next most popular was a creative role — the chef, suggesting recipes and meals, chosen by 40%, and soon after it the DJ, setting the mood with music, at 24%. The next set of answers focused on the sort of emotional support role we’ve begun to see as a pattern: mentor, coach, listener, crisis triage, companion, encourager, partner, friend. Between those supportive and creative options, 80% of respondents picked at least one of those, and 55% picked more than one. They’re all choices that emphasize the emotional role that AIs can play: not just rational tools to make us cleverer, but tools that make us better social humans.

This question turns out to be one of the most revealing in the study, and we will return to it in Section 12, where we break the same data down by segment to show how sharply different kinds of users diverge in the support they want — and how the appetite for relational roles concentrates in one group in particular.

7. How chatbots are rated

Before turning to how people choose between tools, it is worth looking at how they rate the ones they have. Importantly, we only asked each respondent about the chatbots they had told us earlier in the survey that they used: nobody rated a tool they had no experience with. On that

basis we asked how likely they would be to recommend each of their chatbots, on a 0–10 scale, and averaged those answers for each tool.




Figure 6. Average recommendation score (0–10) by chatbot, among the users of each tool (tools with $n \geq 100$ shown).

The leading tools cluster closely together — ChatGPT, Gemini, Claude, and Perplexity all average around 7.0 to 7.3 — while the voice assistants and several of the newer entrants trail behind. The differences between tools are modest, but they conceal much larger differences between the kinds of people doing the rating, which is what the segmentation in Chapter 2 brings out.

8. Paying for chatbots

For all this activity, most of it is free. Across every major tool, the dominant answer to how people pay is that they never have. Only about 30% of respondents pay out of their own pocket for even one chatbot, and just 12% have one paid for by an employer.




Figure 7. Share of each tool's users who pay for it out of their own pocket ($n \geq 100$ shown).

ChatGPT has both the most users and the most paying users, yet even there 70% have never paid a cent and only 17% pay for it themselves. Perplexity (14%) comes next, with the others in single digits. CoPilot is the exception that proves the rule: its users are more likely to have it paid for by an employer (13%) than to buy it themselves (6%), a direct reflection of its place in the workplace. The picture is of a market where free tiers carry the great majority of use, and willingness to pay is real but modest and concentrated in a committed minority — consistent with “avoiding a subscription” ranking as the single most common basis on which people choose a chatbot, which we come to in Chapter 2.

9. Why people walk away

We also asked whether people had ever decided to stop using a particular chatbot, and why. Most could name one. Many of the reasons are ordinary product judgments — another tool gave better answers, or was simpler to work with, or the novelty wore off. But running through the answers is a striking second thread: people reacting not to the product at all, but to the company behind it and how it conducts itself.

Grok attracts the most of this by far, repeatedly abandoned for its association with Elon Musk and X: “I will not use Grok because of its association with X”; “it’s from Elon Musk”; “I distrust the intent behind the answers.” Meta is dropped on plain distrust of the company. Some users were concerned about OpenAI: “I believe that they are being reckless with AI.” Others told us they stopped using DeepSeek over where their data might go. And perceived political bias drives departures from both directions at once — some say ChatGPT “leans way left,” others that it “leans way right” — which suggests the perception is as much about the user’s own lens as the model’s behavior.

The point is that, for a meaningful share of users, the choice of chatbot is partly a values statement: they will walk away over a maker’s ownership, politics, or perceived ethics, more or

less regardless of how good the answers are. That is the same instinct that, among people who don’t use chatbots at all, hardens into never starting — the subject of Chapter 3.

10. If it were a dinner guest

We added two deliberately playful questions. We asked respondents to picture each chatbot as a guest at a dinner party — how would they describe them? — and then to say what animal each one would be. The exercise is light-hearted, and we present it in that spirit. But the answers are revealing precisely because they are unguarded: when people stop rating features and start describing personalities, the emotional relationship they have with these tools comes into view.

Clear and remarkably consistent archetypes emerged, summarized in Table 1.

Chatbot	At the dinner party	As an animal
ChatGPT	The “popular professor” — the authoritative know-it-all with an answer for everything; respected, if sometimes “soulless” or trying too hard to be a friend.	Lion, owl, or loyal golden retriever
Claude	The “thoughtful intellectual” — a careful, creative listener who thinks before speaking.	Owl or cat
Gemini	The “friendly jack-of-all-trades” — the laid-back, always-around helper, ready for anything.	Dog or cat
CoPilot	The “professional workhorse” — all business, focused on getting things done.	Horse or dog
Grok	The “edgy rebel” — witty and unfiltered, sometimes sarcastic or chaotic.	Monkey or hyena
Perplexity	The “precision researcher” — the fact-checker who brings sources and evidence.	Hawk or owl
Siri & Alexa	The “domesticated utilities” — handy for timers, limited in real conversation.	Cat, sloth, or parrot

Table 1. How users described each chatbot as a dinner guest and as an animal.

The individual answers are vivid. ChatGPT is “the smartest one in the room,” “dressed in a suit and tie,” “the old friend that’s been around the longest — everyone can count on them for a favour,” but also “a professor who tries a little too hard to be a friend.” Grok is “a wild teenager on drugs”; Siri, memorably, is “a trained dog that doesn’t hear well.” What stands out is not any single description but the ease and warmth with which people supply them. These are not the answers of people describing a piece of utility software but point to something more conversational, more personal.

Some go further still. One respondent described their main chatbot as “someone who understands my background... they would be exactly like someone who I grew up with, and even share a lot of memories with me.” That is the same emotional register we will meet in the next

chapter among the Emotionally-Led Power Users, who call their favourite tools a “bestie” or “the one that knows me best.” Here, though, it surfaces across the whole population: invited to imagine, people readily cast these tools as characters — wise, loyal, edgy, aloof — with personalities they relate to. The exercise is a joke, but the consistency of the personas, and the affection and irritation woven through them, is a real signal of how emotionally present chatbots have already become in people’s lives.

CHAPTER TWO

Five kinds of chatbot users

How people choose an AI assistant, and the five distinct segments those choices reveal.

11. How people choose a chatbot

So far we have met the everyday chatbot user: tech-comfortable, juggling several assistants, typing more than talking. But averages flatten people. The more interesting question is why someone reaches for one chatbot over another — and there, the population splits apart in revealing ways. We asked our 1,000 respondents what drives their choice of chatbot, across some fourteen dimensions: raw response quality, avoiding subscription fees, style and tone, image and document handling, personalization, integrations, company ethics and policy, and others.




Figure 8. Share of respondents who “always” choose a chatbot on each basis.

Across the whole sample, the loudest signals are practical: most people choose on cost, on the quality of the response, and on the convenience of having one tool that does everything. After those, the ability to interpret and generate different kinds of media matters most. That much is unsurprising. What the averages hide, though, is how differently individual people weigh these factors — and it is precisely those differences, rather than the population-wide ranking, that the next section is built on.

12. Five kinds of chatbot users

We ran a cluster analysis on how people answered the question of how they choose a chatbot. It produced five groups, roughly the same size with one notably smaller, each with a clear personality.




Figure 9. The five segments, by share of chatbot users.

Cluster 1 — Practical Capability Seekers (20%)

These users actively switch tools between work and personal life and choose primarily on capability: response quality, avoiding subscription costs, and media features like interpreting and generating images and documents. They are completely indifferent to emotional connection — the lowest score of any group — as well as to company ethics and corporate policy. Savvy, outcome-focused, treating chatbots as functional instruments.

Cluster 2 — Quality-First Loyalists (21%)

The defining trait is loyalty: they score highest of any group on wanting “one chatbot for everything.” Within that single tool they care deeply about quality and about style and tone, and they’d prefer to avoid paying. Context-switching, integrations, and company policy barely register. They’ve found their assistant and they’re sticking with it.

Cluster 3 — Discerning Multi-Factor Choosers (24%)

The largest and most thoughtful segment. They score highly on every single dimension — something no other cluster does. They split tools across contexts and care about quality, media capability, ethics, company policy, personalization, and integrations all at once. These are power users who have mapped out exactly what they want and demand it across multiple providers.

Cluster 4 — Passive One-Tool Users (23%)

The mirror image of Cluster 3. They want one chatbot for everything and they’d rather not pay — and that’s essentially the whole story. Response quality, personalization, and emotional connection all score low. These are casual users who default to whatever free tool is in front of them and barely think about the choice.

Cluster 5 — Emotionally-Led Power Users (12%)

The smallest segment and, arguably, the most demanding. Every person here said they always choose based on emotional understanding. But they aren't soft, low-effort users — they also score highest across the board on quality, style and tone, media interpretation, and personalization. They want an assistant that performs brilliantly and genuinely gets them. Their language about their favorite tools is telling: “Bestie,” “Companion,” “the one that knows me best.”

The segments live different lives

These aren't just attitudinal differences — the groups differ demographically too.

- **Age and gender.** The Emotionally-Led Power Users are by far the youngest segment (median 46 vs. 56–57 for the two oldest groups), and the most female (63%). The Passive One-Tool Users are the oldest.
- **Work.** Multi-Factor Choosers lead on full-time employment (61%). Loyalists and Passive users have the most retirees. The Emotional segment has the highest student share (13%), consistent with its younger age.
- **Caregiving.** Multi-Factor Choosers and the Emotional segment are the most caregiving-intensive — both show high rates of parenting teenagers and living with elders. Passive users are the most likely to have no caregiving duties and to live alone.
- **Devices.** Multi-Factor Choosers are the most gadget-rich households; Passive users score lowest on every device category. The Emotional segment skews more toward Android than iPhone.

What each segment wants next

Splitting the future-roles question by segment sharpens the divergence. Clusters 1, 2, and 3 overwhelmingly want productive roles — assistant, editor, teacher. The Emotionally-Led Power Users break hard the other way, toward relational roles: listener (41%), best friend (27%), companion (26%), even crisis triage (27%) — every one of these their group's highest score by a wide margin. The Passive users, true to form, lead only on “DJ,” suggesting even their aspirations stay low-effort.

This same group also rates its tools most warmly, and it is here that the recommendation scores from Section 7 come into focus. The Emotionally-Led Power Users give the highest scores of any segment to seven of the ten chatbots we tested — rating ChatGPT at 8.4 out of 10, Gemini at 7.8, CoPilot at 7.5, and Alexa at 7.3. They do not merely use these tools; they bond with them, describing favorites as a “bestie,” a “companion,” or “the one that knows me best.” Their enthusiasm lifts the averages in Figure 6, and helps explain how a tool can post a middling overall score while inspiring genuine devotion in a particular slice of its users. Note that these emotionally-oriented findings come from the smallest cluster (around 117 people), so treat the fine detail as directional rather than precise.

Why segmentation matters

The practical lesson is that there is no single “chatbot user” to design for. A feature that delights the Multi-Factor Choosers — say, deep integrations and document handling — is invisible to the Passive crowd and beside the point for the Emotionally-Led. The clearer the picture of which segment a tool serves, the sharper the product its makers can build. Chatbot designers don’t have to solve every problem for everyone; they can specialize and still be useful.

CHAPTER THREE

The frustration is real

Why a substantial minority don't use chatbots — and the broad distrust of AI and its makers that lies beneath their stated reasons.

13. The people who don't use chatbots

It is tempting to treat people who don't use chatbots as a funnel: a pool of future users who simply haven't gotten around to it yet. That framing is comfortable, and it is mostly wrong. What the non-user study reveals is something more diffuse and more interesting — a broad, somewhat undirected unease about the whole world of AI as it stands in 2026. People are not, for the most part, weighing specific product features and finding them wanting. They are registering a general discomfort: they don't much trust or like the companies involved, and they don't much trust the technology itself.

How we asked

We asked non-users why they hadn't used a chatbot in two different ways, and the order matters. First, an open-ended question:

We noticed you didn't select 'A chatbot, like Claude or ChatGPT'. We're trying to understand how people think about these tools. In your own words, can you tell us what made you decide not to pick this option? For example, was last week an exception; is this a deliberate decision you've made; is it just something you haven't tried yet?

This question just presented a text box: no other framing, no options, just looking for their raw and unfiltered feedback. Then, on a later screen, a closed-ended question: a list of 25 possible reasons, with instructions to check all that applied. Asking the open question first means people had a clean, unprompted chance to say whatever was actually on their minds before any list could steer them.

What they said, cold

Asked first, in their own words, the most common theme by a wide margin wasn't ethics or the environment — it was simply not having a use for the thing. The study's coding of the open-ended responses put “lack of perceived need or general non-use” at the top, well ahead of everything else, followed by deliberate avoidance and a preference for doing one's own thinking.




Figure 10. Themes in non-users' own words, asked before any checklist (n = 101; responses may touch more than one theme).

The free-text answers are quiet and ordinary. “I did not use this last week because I did not need it.” “I never use chatbots or had a reason to.” “For the most part, traditional Google searches fill my information-seeking needs.” Where principled objections do appear, they are often paired with a practical one: “I’m technologically literate, and have no need to use LLMs for any of my daily tasks.” Environmental and ethical concerns are present, but they are not what most people reach for first.

Then came the checklist

Moments later, we handed the same people a list of 25 possible reasons and asked them to check all that applied. The picture changed completely. Non-users now endorsed a great many reasons — an average of 6.2 of the 25 each, with 60% checking five or more — and the ones that rose to the top were overwhelmingly ethical and systemic rather than personal or practical.




Figure 11. Every reason offered on the 25-item checklist, asked after the open-ended question (n = 100). A reminder: this is a small sample, so these figures should be taken as guidance, not absolute numbers.

Taken on its own, this checklist would paint non-users as a principled, AI-skeptical bloc — people kept away by deep concerns about the environment, corporate power, and the ethics of how these systems are built. But we already had their cold answers, and they told a quieter story.

The environmental paradox, and what it really means

Line the two questions up and the contradiction is stark. Environmental impact is the number-one reason on the checklist, selected by 62% — yet only a small share, on the order of 7%, raised it a few screens earlier when asked in their own words. Data and privacy is checked by 53% but volunteered by almost no one. Meanwhile “no real use for it,” the runaway leader when people spoke for themselves at roughly 44%, has no strong equivalent near the top of the checklist.




Figure 12. The same reasons, measured two ways: volunteered unprompted versus selected from a list.

Because we asked the open question first, this cannot be dismissed as fatigue at the end of a long survey. People had a clear, unprompted opportunity to name any reason they liked, and most named a practical one. Only when handed a list of articulate, socially resonant objections did they begin endorsing them in large numbers.

Our read is that the environmental concern, while certainly real, is doing a specific additional job. While the environmental concern is very genuine — people really do worry about the energy and water that AI consumes and the real impacts of noise pollution — the fact that this concern was expressed so much more in the checklist of choices and not in the first, open-ended question suggests strongly that more is going on here. The environmental concern is legitimate and frankly laudable, but it's also the most respectable and most defensible way to voice a much bigger and blurrier feeling. “I don’t trust these companies and I don’t like where this is all going” is hard to fit in a survey box. “I’m worried about the environmental impact” is easy, principled, and admirable — so it soaks up the weight of more nuanced concerns. The numbers are so different precisely because “the environment” is standing in for something larger: a broad distrust of the major companies building AI, and of the technology itself.

Environmental concern is the #1 reason people check off (62%) — yet barely 7% raise it on their own. It is less the reason than the acceptable language for a broader distrust.

What would need to be true

If the barrier is distrust rather than missing capability, the obvious question is what would change it. So we put it to non-users directly, with a final battery: for a long list of possible

conditions, would each make them more likely to use a chatbot, less likely, or neither? The pattern in the answers is the clearest confirmation of the distrust thesis in the whole study.




Figure 13. Conditions that would make non-users more or less likely to use a chatbot (n = 100).

The conditions that move people are almost entirely about trust and restraint, not capability. The strongest draws are a tool that is environmentally responsible (57% more likely), built on a sound ethical foundation (57%), that genuinely respects privacy (53%), that runs locally so that data never leaves the device (51%), that carries no advertising (48%), and that never hallucinates (48%). None of these is a feature in the ordinary sense; every one is a form of reassurance, an answer to a worry. What non-users are asking for is not a cleverer assistant, but evidence that the thing and its makers can be trusted.

The mirror image is just as telling, and it should give pause to anyone building toward an ever-more-present AI. The capabilities that most excite current users are precisely the ones that most repel non-users. Integrating a chatbot into their group chats would make 63% of non-users LESS likely to use one — the single most negative item in the study. Letting it reach into all their apps and tools (39% less likely), having it understand their emotions and feelings (35% less), and going voice-only instead of typing (34% less) all push the same way. More reach, more intimacy, more presence: for this group, that is not the dream but the thing they are backing away from.

63% of non-users say that integrating a chatbot into their group chats would make them LESS likely to use one. More presence is exactly what they don't want.

And then there is the hardest number to design around. Fifty-two percent of non-users agreed with the flat statement that they have no intention of ever using a chatbot or AI if they can avoid it. For half of this group, the honest answer to “what would need to be true?” is “nothing.” That is not a gap to be closed with the right feature or the right campaign; it is a settled position. Taken together, the message is consistent: what little would move non-users is trust, not capability — and for many, not even that. Pushing more AI deeper into their lives is the surest way to push them further away.

How many people use these tools

It helps to put non-use in context, because adoption estimates vary widely depending on who is counting. In our own studies, about 77% of entrants to the main survey had used a chatbot in the past week, and screening in the non-user study implied a similar figure of roughly 76%. A March 2026 NBC News / Hart Research poll of 1,000 registered voters found 56% had used an AI platform such as ChatGPT, Copilot, or Gemini in the past two or three months, and a widely-cited mid-2025 industry estimate put combined weekly active use across the major assistants at around 64%. Our numbers run high, almost certainly because a Prolific panel skews more online than the country as a whole. But the more useful reading is the flip side: even on the most generous count, a substantial share of people consistently opt out — and many of them, as the answers above show, have decided this is not for them rather than merely delayed.

Pull the threads together and the picture is consistent: the distrust is real, and much of it is aimed less at any particular product than at AI itself and the companies behind it. “I just don’t like AI,” one non-user wrote. “It makes things up and can’t be trusted.” Others framed it as resistance to something being imposed on them — “a fad that is being shoved down our throats by the technology companies that have so heavily invested in them” — or as a settled wariness of corporate motives: “I don’t trust... private companies not to take full advantage of any way to make money or take advantage of people in general.”

When people do name names, the targets are familiar and the language is unsparing. Meta draws a lot of heat — “I don’t trust it, so I don’t use it anymore”; “it’s just collecting your info rather than being actually helpful”; “I don’t like Meta as a company.” So does Google — “I no longer trust Google to safeguard my data,” one user explained, citing its willingness “to share users’ information” — and Microsoft, whose Copilot one respondent dismissed as hopelessly “tangled with all of Microsoft’s other offerings.” Even the market leader isn’t exempt: “I don’t trust ChatGPT to keep my information private.”

Often the distrust attaches to the people at the top as much as the products. Grok is the clearest case: people abandon it “because of its association with X,” because “it’s from Elon Musk,” and because “I distrust the intent behind the answers.” The judgments get personal and harsh — “Grok, I don’t trust it or Musk,” and, most bluntly of all, “I stopped using Meta and Grok because their CEOs are full-blown sociopaths.” This is no longer a conversation about features; it is a verdict on the founders.

The throughline is that, fairly or not, a chatbot arrives carrying the reputation of the company — and the figurehead — that builds it. For this group, no feature ships without that baggage, and that, more than any missing capability, is what keeps them at a distance.

CHAPTER FOUR

Conclusions

What the two studies add up to.

Two ideas run through this work. The first is that there are real opportunities in focusing on specific user needs. There is no single user to win, and trying to serve everyone serves no one well — but a tool built around the needs of a particular group can be genuinely valuable to the people in it. The second is that the better anyone can anticipate how people will respond before acting, the stronger their decisions will be. Research like this replaces a guess with a portrait.

On the user side, that portrait is encouraging. Chatbot use is a plural, healthy ecosystem in which people already mix several tools, reach for them across phone and computer, and — strikingly — describe them not just as software but as characters, the popular professor or the edgy rebel, with personalities they relate to. Their use is mostly mundane and practical, but the appetite for emotional and relational roles is real and concentrated, and it points to room for tools that do more than answer questions.

The non-user side is harder, and it is where this research most changes the picture. The instinct is to treat non-use as a problem of capability or persuasion — a better feature, a clearer pitch. The data says otherwise. What keeps these people away is not a gap in what chatbots can do but a judgment about who is making them. The distrust attaches to the companies and, increasingly, to the founders themselves; a chatbot arrives carrying the reputation of its maker, and no feature ships without that baggage. That cannot be closed with a capability or a campaign. Trust of that kind is earned slowly, through conduct, and it can be spent across a whole industry by the actions of a few of its most visible players.

This has an uncomfortable corollary worth stating plainly: the moves that most delight current users — deeper integration, more emotional fluency, a more constant presence — are precisely the ones non-users recoil from. Winning them is not a matter of bringing more AI more insistently into their lives; for many, that is the very thing they are backing away from, and for some the honest answer is that nothing would change their minds soon. The most valuable thing a maker can take from this is not a feature roadmap but a sense of proportion: clarity about which users it is actually for, humility about how much of the resistance is about trust rather than technology, and the patience to understand that distrust before trying to dissolve it.