

# Your Second Brain Has a Landlord

*Your AI remembers everything about you, and someone else owns the memory.*

**Kymata Labs Research** · An independent research institution studying how AI systems actually behave in-12 min read production.

Tags · AI Privacy · Data Governance · Persistent Memory · Digital Rights

As AI assistants gain persistent memory, they are becoming the most intimate dossier ever assembled on a person, your fears, health, finances, relationships, and 2 a.m. questions, and you do not own it, cannot fully delete it, and cannot take it with you. This paper traces the 2024–2026 rollout of memory across ChatGPT, Gemini, Claude, and Microsoft 365 Copilot; the gap between deleting a chat and erasing its memory; the court order that overrode users’ deletion rights for six months; the sixtyfold extension of one retention window; and survey evidence that most people do not know their data trains models or whether they can opt out. The celebrated “second brain” is rented, and the lease terms are written by whoever runs the model.

## The argument, in moves

1. **Memory turned the assistant into a dossier.** Persistent memory means the system no longer answers a question and forgets it. It accumulates a model of you across every session.
2. **You do not own the derived profile.** You can export a transcript. You cannot export, port, or fully see the inferences the system has drawn on top of it.
3. **Deletion is two switches, and you only touch one.** Deleting a chat does not delete the memory derived from it. The derived profile persists unless you remove it deliberately.
4. **Your deletion settings can be overridden.** A third party’s lawsuit forced one provider to retain “deleted” chats for roughly six months. The setting did what it promised; a court overruled it.
5. **Consent is real on paper and absent in practice.** Most people do not know their data can train a model or whether they can opt out, and defaults increasingly arrive switched on.

The one-liner: the second brain is the most valuable personal database ever built, and you are the tenant, not the owner.

## Memory arrived everywhere, at once

In the span of about eighteen months, persistent memory went from a novelty to a default across every major assistant. The dates matter, because they show how fast the most intimate data surface in computing was switched on.

Assistant	Memory milestone	When
-----------	------------------	------

ChatGPT	Memory announced; GA; then references all past chats	Feb 2024 → Apr 2025
Gemini	Personal Context, on by default (opt-out)	Aug 2025
Claude	Memory to Team/Enterprise, then Pro/Max	Sep–Oct 2025
Microsoft 365 Copilot	Copilot Memory general availability	Jul 2025

Sources: OpenAI, Google, Anthropic, Microsoft product announcements (2024–2025).

The pattern is uniform: memory ships, then expands, then becomes the default. Gemini’s Personal Context arrived opt-out, which means the private state, no long-term memory, was the one you had to go and choose.

## The most intimate dossier ever assembled

---

Think about what a long-running assistant actually accumulates. Not a browsing history, not a purchase log, but the unedited interior: the health worry you typed at midnight, the draft of the resignation letter, the relationship you are unsure about, the finances you would not show a friend. People talk to these systems with a candor they reserve for almost no one, precisely because it does not feel like talking to a company. It is. Every disclosure is a row in a database owned by someone else, and memory is the feature that makes the rows persistent and joinable into a portrait.

## Deletion is not erasure

---

Here is the gap most people never see. Deleting a conversation does not delete the memory derived from it. In OpenAI’s own framing, deleting a chat does not erase its memories; you must delete the memory itself. Those are two separate controls, and almost everyone only ever uses the first. The transcript disappears; the conclusions the system drew about you, the derived profile, remain until you go and remove them deliberately. The interface makes the easy action feel complete when it is not.

## The preservation order: your deletion, overruled

---

Even the deletion that does work can be suspended by forces you never agreed to. In the New York Times litigation against OpenAI (No. 1:23-cv-11195, S.D.N.Y., before Magistrate Judge Ona T. Wang), a preservation order entered on May 13 2025 required OpenAI to retain output log data that would normally be purged within 30 days, including deleted chats, Temporary Chats, and API content, across ChatGPT Free, Plus, Pro and Team and the API. The obligation was lifted on September 26 2025 <sup>5</sup>. For roughly six months, your deletion settings did exactly what they promised, and a third party’s lawsuit overrode them anyway. The lesson is not that OpenAI behaved badly; it is that your control over your own data is contingent on litigation you are not a party to.

## Retention, quietly extended

---

The window itself can move. On August 28 2025, Anthropic updated its consumer terms: consumer chats and Claude Code sessions are used to train models when the setting is on, and data retention was extended from 30 days to five years, with a decision deadline of October 8 2025 <sup>7</sup>. A thirty-day memory and a five-year memory are different products with different risk profiles, and the change arrived as a terms update with an opt-out clock, not a feature you chose.

## Consent that nobody read

---

All of this is, technically, in the terms you agreed to, and that is precisely the problem. In a representative survey of 3,270 UK adults, 53.8% did not know their data could be used to train a model, and 76.5% were not sure whether they could opt out <sup>6</sup>. Consent that the consenting party does not understand is consent in name only. The defaults, the deletion gaps, and the retention clocks all operate inside that gap of understanding.

## No way out: the second brain does not port

---

Suppose you want to leave. You can export your conversation history, the transcript of what you typed and what the model replied. What you cannot do is export the derived memory profile or move your memory to another provider. There is no portability standard for a second brain. The transcript is yours to download; the model of you, assembled from those transcripts and the inferences layered on top, stays with whoever runs the model. Leaving means starting over from zero, which is exactly the switching cost that makes the landlord a landlord.

## The lease

---

A second brain is a genuinely useful thing. The argument here is not to refuse it; it is to read the lease before you move in your whole interior life. Treat every assistant's memory settings as lease terms: what is retained, for how long, whether it trains the model, and how to delete the derived profile rather than just the chat. Know which defaults arrived switched on. And recognize that even good settings sit downstream of corporate policy changes and court orders you do not control. The data is the most personal you will ever generate. Decide, deliberately, what the landlord gets to keep.

## Frequently asked questions

---

### **Can't I just delete my data?**

You can delete conversations, and you should be able to delete stored memories. But deleting a chat does not delete the memory derived from it. In OpenAI's own words, "deleting a chat doesn't erase its memories; you must delete the memory itself." Those are two separate switches, and most people only ever touch the first. The derived profile, what the system has concluded about you, persists unless you remove it deliberately.

### **Can't I export my data and move to another provider?**

You can export your conversation history, the transcript. What you cannot do is export or port the derived memory profile across providers. The transcript is yours to download. The model of you, assembled from those transcripts and the inferences on top, stays with whoever runs the model. No portability standard exists for a second brain, so leaving means starting over.

### **If I never turn memory on, am I safe?**

Partly, and pay attention to the defaults. Google's Gemini Personal Context arrived on by default in August 2025 (opt-out rather than opt-in), which means the safe state was the one you had to go and choose. Read what is on before you assume it is off.

### **Surely a deleted chat is gone in 30 days?**

Usually, but "usually" is doing a lot of work. For roughly six months in 2025, a court preservation order in the New York Times litigation required OpenAI to retain output log data that would normally be purged within 30 days, including deleted and Temporary Chats and API content. The order was entered May 13

2025 and lifted September 26 2025. Your deletion settings did what they promised, and a third party's lawsuit overrode them anyway.

### **Isn't all of this in the terms I agreed to?**

It is, and that is the problem, because almost nobody knows what they agreed to. In a representative survey of 3,270 UK adults, 53.8% did not know their data could be used to train a model, and 76.5% were not sure whether they could opt out. Consent that nobody understands is consent in name only.

### **What's the single most useful thing I can do today?**

Open the memory settings on every assistant you use and read them as a lease: what is retained, for how long, whether it trains the model, and how to actually delete the derived profile rather than just the chat. Then decide, deliberately, what this landlord gets to keep. The defaults were not written for you.

## References

---

- 1 OpenAI (2024). "Memory and new controls for ChatGPT." Announced Feb 13 2024; GA Sep 2024; expanded Apr 2025 to reference all past chats. Includes the statement that deleting a chat does not erase its memories. <https://openai.com/index/memory-and-new-controls-for-chatgpt/>
- 2 Google (2025). "Temporary chats and privacy controls in the Gemini app." Personal Context introduced Aug 13 2025, on by default (opt-out). <https://blog.google/products-and-platforms/products/gemini/temporary-chats-privacy-controls/>
- 3 Anthropic (2025). "Using memory with Claude." Memory reached Team and Enterprise around Sep 2025, then Pro and Max in October 2025. <https://support.claude.com/en/articles/11817273>
- 4 Microsoft (2025). "Introducing Copilot Memory." Microsoft 365 Copilot Memory reached general availability July 2025. <https://techcommunity.microsoft.com/blog/microsoft365copilotblog/introducing-copilot-memory-a-more-productive-and-personalized-ai-for-the-way-you/4432059>
- 5 OpenAI (2025). "Response to The New York Times' data demands." On the May 13 2025 preservation order in NYT Co. v. OpenAI, No. 1:23-cv-11195 (S.D.N.Y.), before Magistrate Judge Ona T. Wang; obligation lifted Sep 26 2025. <https://openai.com/index/response-to-nyt-data-demands/>
- 6 Grosse, M. & Ebert, T. (2025). "What do people know about how their data is used to train AI?" Representative survey of 3,270 UK adults. arXiv:2510.27275. <https://arxiv.org/html/2510.27275v1>
- 7 Anthropic (2025). "Updates to Consumer Terms and Privacy Policy." Aug 28 2025: consumer chats and Claude Code sessions used for training when the setting is on; retention extended from 30 days to five years; decision deadline Oct 8 2025. <https://www.anthropic.com/news/updates-to-our-consumer-terms>