

Public AI tools at work: what providers say happens to the text staff paste in

Someone on your team has probably pasted a customer email, a contract clause or a block of code into a public AI tool. Whether that text can be used to train the tool depends on the plan the account is on. We read what OpenAI, Anthropic, Google and Microsoft publish about it on 28 October 2026.

Author	XONIK Editorial
Published	October 27, 2026
As of	October 28, 2026
Reading time	8 min read
Online	https://xonik.com/insights/market-intelligence/public-ai-tools-at-work-what-providers-say-about-your-data/

IN BRIEF

Whether text pasted into an AI tool can be used to train it depends on the plan, not the brand. OpenAI, Anthropic and Google say individual accounts may be used for training or review unless the user opts out, and business plans are not used by default. Microsoft says Copilot prompts and responses are not used to train foundation models.

KEY FINDINGS

1. The plan on the account sets the default.

OpenAI says content from Free, Plus and Pro may be used to train its models unless the user opts out, and that Business, Enterprise, Edu and the API are not used by default. (Source: OpenAI Help Center, How your data is used to improve model performance, as of 28 October 2026.)

2. A feedback button can override a setting.

OpenAI says the whole conversation behind a thumbs up or down on an individual plan may be used even after the user opted out. Anthropic says feedback or bug reports on its commercial products may be used to train. (Sources: OpenAI Help Center; Anthropic Privacy Center, as of 28 October 2026.)

3. Some individual accounts involve human reviewers.

Google says trained reviewers, including service providers, read some Gemini app chats, and asks people not to enter confidential information they would not want a reviewer to see. For Workspace it says content is not human reviewed outside your domain without permission. (Sources: Google, Gemini Apps Privacy Hub; Google Workspace Help, as of 28 October 2026.)

4. Business tools work inside your permissions.

Microsoft says Copilot only surfaces organizational data the user already has permission to view, and that prompts and responses are not used to train its foundation models. Google describes the same access rule for Gemini in Workspace. (Sources: Microsoft Learn; Google Workspace Help, as of 28 October 2026.)

5. Temporary and incognito modes sit outside training.

OpenAI says temporary chats are not used to improve its models while they stay temporary. Anthropic says Incognito chats are not used to improve Claude. (Sources: OpenAI Help Center; Anthropic Privacy Center, as of 28 October 2026.)

What does OpenAI say?

OpenAI separates individual plans from business ones. For Free, Plus and Pro, it says it may use your content to train its models, and that you can opt out by turning off a setting called Improve the model for everyone, or through its Privacy Portal. Temporary chats are not used while they stay temporary. If someone saves one, it becomes a regular chat and follows the account's settings.

For ChatGPT Business, Enterprise and Edu, and for the API, OpenAI says it does not use inputs or outputs to improve its models by default. Its enterprise privacy page adds that business data may pass through automated classifiers and safety tools, and that human review depends on the service. The catch on individual plans is feedback. OpenAI says the entire conversation attached to a thumbs up or down may be used to improve its models even after you opted out. On business plans it says feedback does not permit training.

What does Anthropic say?

Anthropic's page for its consumer plans, Free, Pro and Max, names three cases in which chats and coding sessions may be used to improve Claude: you choose to allow it, a conversation is flagged for safety review, or you opt into a program such as its trusted tester program. It says Incognito chats are not used to improve Claude, even when model improvement is switched on.

For commercial products such as Claude for Work and the API, Anthropic says it does not use inputs or outputs to train its models by default. The exception is feedback: if you explicitly report feedback or bugs, for example with the thumbs button, it may use those chats and coding sessions to train. Owners of Team and Enterprise plans can switch the feedback button off for their organization.

What does Google say?

Google's help page for the consumer Gemini app says that, with Keep Activity on, your activity can be used to improve Google services, including the generative AI models behind them. It says trained reviewers, including service providers, read some chats. With the setting off, it says future chats are not used to improve Google's models unless you send feedback.

Workspace is a separate product with separate terms. Google's privacy hub for Gemini in Workspace says your content is not human reviewed or otherwise used for generative AI model training outside your domain

without permission, and that Gemini only reaches content the user can already open. Administrators set how long prompts and responses are kept.

What does Microsoft say?

Microsoft's page on Copilot at work says prompts, responses and data accessed through Microsoft Graph are not used to train the foundation models behind Copilot. It says Copilot only surfaces organizational data the user already has permission to view, so the permissions already in your tenant decide what it can reach. It also says interactions are stored, and that administrators can search them and set how long they are kept.

What each provider says about training, by type of account, as of 28 October 2026

Provider	Account type	What the provider says	Source and date
OpenAI	Free, Plus, Pro	Content may be used to train models; the user can opt out; temporary chats are not used while temporary; feedback can bring a whole conversation back in	OpenAI Help Center, How your data is used to improve model performance (relative update time only); as of 28 October 2026
OpenAI	Business, Enterprise, Edu, API	Not used to improve models by default; feedback does not permit training	The same page; OpenAI, Enterprise privacy at OpenAI, updated January 8, 2026
Anthropic	Free, Pro, Max	May be used if the user chooses, if a chat is flagged for safety review, or if the user joins a program; Incognito chats are not used	Anthropic Privacy Center, last updated March 16, 2026
Anthropic	Claude for Work, API	Not used to train by default; feedback or bug reports may be used	Anthropic Privacy Center, no date shown; as of 28 October 2026
Google	Gemini app, personal account	With Keep Activity on, activity can improve Google's models; trained reviewers read some chats; do not enter confidential information	Google, Gemini Apps Privacy Hub, updated October 23, 2026
Google	Gemini in Workspace	Not human reviewed or used for model training outside your domain without permission	Google Workspace Help, Generative AI in Google Workspace Privacy Hub, last updated October 6, 2026
Microsoft	Copilot at work	Prompts, responses and Graph data are not used to train foundation models	Microsoft Learn, Data, Privacy, and Security for Microsoft Copilot; as of 28 October 2026

The brand is the same in each pair of rows. The default is not.

What does this mean for a staff rule?

A rule that names brands goes out of date and misses the point. The same person can use ChatGPT on a personal account in the morning and in a company workspace in the afternoon, and the documents give those two accounts different defaults. The line to draw runs between accounts the company administers and accounts it does not.

On a personal plan, the protection is a setting each person has to find and keep switched off, with a feedback button that can undo it for one conversation. On a business plan, the default is set for everyone and an administrator holds the controls. A rule should write that difference down: which accounts are approved, who administers them, and what may never go into the others.

In the patterns we see, approving a business-plan tool that people actually want to use works better than banning the personal ones, because a ban moves the use out of sight. That is our reading, and the choice is yours. What you cannot do is learn the answer from a provider's page, because the pages describe the products and say nothing about your staff's habits.

When is a public AI tool acceptable for a task?

A starting rule to adapt, built on the documents above. It is not legal advice, and a signed agreement with a provider overrides a public page.

Reasonable on an account the company administers

- Drafting and editing text that is already public or about to be published
- Working on internal material the person is allowed to use, where the business plan's default of no training applies
- Summarizing a document the person has the right to share beyond their own team

Keep out of any account the company does not administer

- Customer or employee personal data
- Contract terms and prices under negotiation
- Unreleased financial results
- Credentials, keys and source code that is not public
- Anything covered by a confidentiality agreement with a third party

A stricter rule protects more and gets used less. A looser rule gets used more and protects less. Pick the risk you could explain to a customer.

The brand tells you little. The plan on the account tells you what happens to the text.

(Sources: OpenAI Help Center; Anthropic Privacy Center; Google, Gemini Apps Privacy Hub; Google Workspace Help; Microsoft Learn, as of 28 October 2026.)

What should you watch for in these documents?

These pages change without notice, and some carry no date. OpenAI's help article showed only a relative update time when we read it. Settings can also sit apart from the account setting: OpenAI says its coding agent Codex has a separate setting for training on full environments, and that changing the ChatGPT setting does not change it. Product names move as well. Microsoft's page notes that Microsoft 365 Copilot is now named Microsoft Copilot, so a rule that lists product names needs an owner who rereads them.

Method, as-of date and limits

This briefing describes the position on 28 October 2026. We opened each provider page that day and read it in full: OpenAI's help article and enterprise privacy page, Anthropic's two privacy center articles, Google's Gemini Apps help page and its Workspace privacy hub, and Microsoft's Copilot privacy page. Statements are attributed to the provider that made them. Where we give a reading of our own, we say so.

- Four providers are covered. Others, including Meta, Mistral and many smaller tools, are not.
- The pages are the providers' own descriptions. They do not show how a product behaves in practice, and we did not test any product.
- Several pages show no date or only a relative one, so they may have changed since we read them.
- Plan and product names change, and a plan not named here is not covered.
- A signed agreement with a provider governs your organization. This briefing is not legal advice.

Sources

We read each of these pages on 28 October 2026.

1. [OpenAI Help Center, How your data is used to improve model performance, relative update time only: training defaults by plan, opt-out, temporary chats and feedback](#)
2. [OpenAI, Enterprise privacy at OpenAI, updated January 8, 2026: business data, ownership and review](#)
3. [Anthropic Privacy Center, Is my data used for model training? \(consumer plans\), last updated March 16, 2026](#)
4. [Anthropic Privacy Center, Is my data used for model training? \(commercial products\), no date shown](#)
5. [Google, Gemini Apps Privacy Hub, updated October 23, 2026](#)
6. [Google Workspace Help, Generative AI in Google Workspace Privacy Hub, last updated October 6, 2026](#)
7. [Microsoft Learn, Data, Privacy, and Security for Microsoft Copilot](#)

USE THIS

[Check who answers for the AI you use](#)

[Start from the AI governance policy template](#)

List the AI tools your teams use this week and write beside each one whether the account is personal or administered.

This document was made from the article at

<https://xonik.com/insights/market-intelligence/public-ai-tools-at-work-what-providers-say-about-your-data/>. The online article is the current version.