PD **PRODUCTS DECODED**
Operator notes on AI, PM, and the stuff nobody tells you.

GUIDE · PRICES CHECKED 7 OCT 2026

Claude Code vs Cursor vs Codex

What Your $20 Actually Buys (Oct 2026)

All three AI coding tools start at $20 a month. Where they run, who trains on your code, where the limits bite, and three alternatives, with every price sourced from the official pages.

productsdecoded.com

By Debashis Nayak

You've probably seen all three of these in your feed this month, usually in a thread arguing that one of them is the only one worth paying for. They all write code. They all start at $20 a month. They are built for quite different ways of working, and the differences that cost you money or data sit in the pricing footnotes.

Quick bias check before you read on. I build most of my own products with Claude Code, in a terminal, most days. For Cursor and Codex I haven't put in the same hours, so everything below about them comes from their own pricing and documentation pages. I read all of those on 7 October 2026, and the links are at the bottom.

Which one fits you, in ten seconds

- **You live in a terminal and want an agent that works through your repo:** Claude Code.
- **You want an editor first, with AI in every keystroke and a choice of models:** Cursor.
- **You already pay for ChatGPT Plus and want to hand off tasks that run in the cloud:** Codex.

That's the short version. The rest is the fine print that decides which one you'll still be paying for in three months.

Claude Code: an agent in your terminal

Claude Code is Anthropic's coding agent. It reads your codebase, edits files and runs commands. The terminal is the core, and the same engine also runs as a VS Code and JetBrains extension, a desktop app, on the web and in the Claude mobile app. Your project's `CLAUDE.md` files, settings and MCP servers carry across all of those, which is the part I'd miss most if I switched.

It runs on your machine by default. Prompts and outputs go to the model over TLS. Sessions started from the web or your phone run in isolated machines that Anthropic manages.

It's good at multi-file features and bug fixes, the boring stuff like tests, lint, merge conflicts and dependency bumps, and anything you want to script: `claude -p` runs it non-interactively, so it fits in CI.

The catch: there's no free tier, and the limits are hard. Usage resets on a five-hour window plus a weekly cap. Switching models doesn't get you more, and going past the cap means buying usage credits.

Cursor: an editor that lets you pick the model

Cursor is a code editor built around AI, plus a CLI and cloud agents you can start from Slack, the web or your phone. Its standout is model choice: the docs list frontier models from OpenAI, Anthropic, Google, SpaceXAI and more. Paid plans also get unlimited Tab completions.

The editor runs locally, but every request goes through Cursor's servers, even if you bring your own API key. Cursor says that's where it does its final prompt building. Cloud agents run in isolated machines.

The catch: the $20 doesn't cover heavy agent use. Cursor's own docs say daily agent users typically spend $60 to $100 a month in total, and power users often $200 or more. Third-party models draw from the included pool at API prices.

Codex: the coding agent inside ChatGPT

Codex is OpenAI's coding agent, and it comes with a ChatGPT subscription, so there's no separate product to buy. It works two ways. The CLI and IDE extension work against your local repo, with a sandbox you control. Codex Cloud gives each task its own workspace, and those tasks keep running while your computer sleeps.

It's good at tasks you can delegate and walk away from, terminal and CI work (`codex exec`), local code review (`/review`), and GitHub PR review by mentioning `@Codex` .

The catch: usage is shared and capped. OpenAI's pricing page says Codex shares usage with ChatGPT Work, and cloud tasks burn through it faster. OpenAI publishes message limits as estimated ranges per five hours, not fixed numbers, so I won't quote one here.

Pricing, every tier

	Claude Code	Cursor	Codex
Free	Not included on Claude's Free plan	Hobby, with limited agent requests	Free $0 and Go $8/mo, desktop app only, "subject to rollout"
Entry	Pro $20/mo ($17/mo billed yearly)	Pro $20/mo	Plus $20/mo
Middle	Max 5x $100/mo	Pro Plus $60/mo	Pro from $100/mo
Top	Max 20x $200/mo	Ultra $200/mo	Pro at $200 or $500/mo
Teams	$20/seat/mo yearly ($25 monthly); premium seat $100	$40/user/mo; premium $120	Business $20/user/mo yearly ($25 monthly), 2+ users

Prices are in USD before tax. If you're in India, Cursor has a "Start" plan at ₹649 a month including tax, but it only covers Cursor's own models. Codex's Free and Go plans only run in the desktop app. Plus is the cheapest tier that adds the web, CLI, IDE and code review.

Who trains on your code

This is the table I'd read twice before connecting a work repo.

	Personal plans	Team and business plans
Claude Code	You choose: the pricing page lists model training as opt-out, and Anthropic trains on your sessions only when the setting is on	Not used for training by default
Cursor	Only if Privacy Mode is off. Privacy Mode is free on every plan	Privacy Mode is on by default for Enterprise
Codex	Used for training unless you opt out ("Improve the model for everyone"). Codex has a separate "Include environments" setting too	Not used for training by default (Business, Enterprise, Edu, API)

On retention: Anthropic keeps consumer data for 30 days with training off and five years with it on. Cursor's Privacy Mode comes with zero-data-retention agreements with its model providers, with a few models as exceptions that need admin approval. I found no official statement of zero data retention for Codex, so assume it doesn't have one until OpenAI says otherwise.

Where the $20 runs out

All three give you the same entry price, and all three meter you differently after that.

- **Claude Code** cuts you off. Five-hour windows plus a weekly cap, shared across models. You wait, upgrade to Max, or buy credits.
- **Cursor** keeps going, and the bill grows. Agent use on third-party models comes out of an included pool at API prices, which is how daily users end up at $60 to $100.
- **Codex** shares one allowance with ChatGPT Work, so heavy use there eats into your coding, and cloud tasks drain it faster than local ones.

If you code with an agent every day, budget for the middle tier of whichever you pick. The $20 plan is a trial for that kind of use.

Three alternatives worth a look

Tool	What it is	Price
GitHub Copilot	GitHub's assistant: chat, agent mode, code review, a cloud agent and a CLI	Free · Pro $10/mo · Pro+ $39 · Max $100
Gemini CLI	Google's terminal agent	Free for 1,000 requests a day with a Google account
Devin Desktop (formerly Windsurf)	An IDE with a command centre for local and cloud agents. Cognition renamed Windsurf on 2 June 2026	Free · Pro $20/mo · Max $200

Two notes. Since 24 April, GitHub may train on interactions from Copilot Free, Pro and Pro+ users, code snippets included, unless they opt out. And Google's Gemini CLI quota page was last updated on 18 June 2026, so check it before you rely on that 1,000 number.

Sources

All fetched on 7 October 2026 unless a date is given.

- Claude Code: overview, data usage, costs, Claude pricing, Max plan
- Cursor: pricing, models and pricing, data use (dated 3 Sep 2026), privacy and governance, security (dated 25 Aug 2026)
- Codex: pricing, CLI, cloud, how OpenAI uses your data
- Alternatives: GitHub Copilot plans, Gemini CLI quotas (updated 18 Jun 2026), Devin pricing, Windsurf is now Devin Desktop

Pricing pages for AI tools change every few weeks, so some of this will be out of date before the year ends. If you spot a number that's moved, tell me and I'll update the page.