Gemini 4 Carbon Ships at a Fraction of Fable 5.2's Cost as Anthropic Starts Its IPO Roadshow

AI High Signal Digest

2026-11-11

Gemini 4 Carbon Ships at a Fraction of Fable 5.2's Cost as Anthropic Starts Its IPO Roadshow

By AI High Signal Digest • November 11, 2026

Google's strongest model is now generally available and comes in just behind Claude Fable 5.2 at much lower cost per task. Also covered: Anthropic's IPO pricing, faster AI-timeline forecasts after a METR estimate, Claude's new work tools, and a run of robotics releases.

Gemini 4 Carbon is generally available, and the case for it is cost

Google opened Gemini 4 Carbon to developers, enterprises and consumers two weeks after announcing it. Developers can use it through the paid Gemini API, and Google AI Ultra subscribers get it in AI Studio and Antigravity [1]. Through November 30 it costs the same as Argon: $2 per 1M input tokens and $10 per 1M output tokens. From December 1 that rises to $4/$20, with cached input 95% off. Both models share a 1M-token context window, so switching between them means changing one model string [1]. Google positions Carbon for the hardest multi-step engineering, formal math and vulnerability research, and Argon for most everyday coding and knowledge work [1]. Carbon was trained entirely on Google TPUs [2].

Independent results put Carbon just behind Claude Fable 5.2 on quality, at a much lower cost per task:

- **Vals Index:** #3, about 1 point behind Fable 5.2, at $10.51 per test versus $17.94. It stays #1 on Vals' Finance Index, Legal Research and Tax [3].
- **MKA:** #2, 3 points behind Fable 5.2, at $6.29 per task versus $31.97. MKA is a new Google DeepMind / Artificial Analysis benchmark of 200 private, expert-written tasks on visually complex files. Each task takes a professional 3–8 hours [4].

- **Perplexity WANDR:** 0.857 at $5.86 per task, which Perplexity says matches Fable 5.2 at 40% of the cost [5].
- **CursorBench:** top model at 60.4% at high effort [6].

The low cost does not hold everywhere. On Roboflow's vision evaluations at high effort, Carbon was the slowest model tested at 81.06 seconds per image, versus 9.91s for Opus 5.5 [7]. It was also the most expensive at $91.89 per 1,000 images [8].

The biggest technical result tied to the launch is a DeepMind team's proof of key security properties of the machine code behind the Signal protocol. The team says this is the first time production code of a widely deployed protocol has been formally verified at this depth, by any method. The work took a month, against what it says could be decades of combined effort for similar projects [9]. It used AlphaProof Code, an agentic framework running on LLMs including Carbon, and the specifications and proof artifacts are being open-sourced [9].

Fable 5.2: stronger on real work, but token-hungry

More data arrived on the cost problem flagged when Fable 5.2 launched. SemiAnalysis says that on its materials-science work tasks, Fable 5.2 costs 89% more than GPT-6 Astra and uses 4.3× the output tokens for about the same score. Against Carbon the cost gap is 7.5×. SemiAnalysis blames the model's persistence and its habit of spawning subagents [10]. On software-engineering tasks, Fable 5.2 nearly triples Opus 5.5's score, but cost per task rises by about the same multiple [11].

On the Remote Labor Index, which uses real paid freelance projects, Fable 5.2 completed 41.3% of projects to at least the standard of the human original. Fable 5.1 managed 17.9% [12]. CAIS says the best AI fully automated 2.5% of these projects when the index launched a year ago, and that the rate is doubling about every three months [13]. On Vals' new RSI Index, which tests AI-research tasks, the gain is small: Fable 5.2 scores 37.8, against 37.3 for Opus 5.5 [14]. Vals notes that Anthropic says it deliberately held back training gains in areas like kernel optimization, so results from released models may understate what Anthropic has internally [15].

Anthropic sets IPO terms below expectations

Anthropic launched the roadshow for its proposed IPO. Its SEC registration statement is filed but not yet effective [16]. A widely shared summary reports 825 million shares at $70–$78 each. That implies a fully diluted valuation of about $1.24–$1.38 trillion, below the roughly $1.5 trillion some had expected [17].

Timeline forecasts pulled into 2027

The AI Futures Project now expects AI that matches top human AI researchers to likely arrive in 2027, barring a major slowdown, and says progress is five months ahead of its AI 2027 scenario. The update follows a METR estimate that Anthropic's progress has roughly doubled in speed, with AI driving much of the speedup [18]. Eli Lifland moved his median for a superhuman AI researcher from about 2030 to 2027 [19]. Vals' trend line for Anthropic's released models points to August 2027 for fully automated AI research [15]. Not everyone agrees: Tamay Besiroglu predicts the intelligence explosion will not happen and that the people forecasting it will be discredited [20].

Claude adds office tools

Anthropic is rolling out a set of work tools in Claude [21]:

- **Meetings** (Team and Enterprise): Claude joins Google Meet calls and starts follow-up work during the call.
- **Sheets:** builds a spreadsheet, with formulas, from a description.
- **Whiteboards:** turns a team brainstorm into a shared board with deadlines.
- **Data:** now out of beta.
- **Connectors:** work inside Adobe, Atlassian and Box.

Separately, MCP Apps can now serve a full site inside Claude and give Claude tools to act on what is on screen [22]. That feature is in public preview and details may change [23].

Robotics releases

- **Gemini Robotics 2.5 (Google DeepMind):** learns hundreds of everyday skills [24]. Gemini 4 Argon helped build it by generating virtual practice spaces from a single camera view [25]. In a proof of concept with Apptronik and Mercedes-Benz, an Apollo 2 robot picked and organized car parts with 96% accuracy [26].
- **GR00T 2 (NVIDIA):** an open model that runs entirely on the robot's own compute. It targets tasks of 20+ steps lasting 5+ minutes. Training data (100K hours) and tooling are available now; the weights are due in December [27].
- **LingBot-Touch (Robbyant, part of Ant Group):** an open-source 5.46M-parameter touch encoder. Adding it to Physical Intelligence's 0.5 raised success on 26 real-world tasks from 48% to 70% [28].

Also notable

- **MiniMax H3-Image:** an open-weight model that generates 16MP images in about 10 seconds and accepts up to 16 reference images [29]. One

reader of the license says it includes a $1M revenue gate and bars use in the EU, UK, South Korea and US [30].

- **Goodfire interpretability:** Goodfire argues that neurons are best understood by their computational function rather than by whether they are polysemantic [31]. It reports that ablating the weakest 92% of activations has only modest effects on capability [32].
- **Qwen4:** Alibaba put Qwen4-Max into preview [33]. Early user reports are poor. One test scored it below Qwen 3.8 Flash [34], and another tester said it did not feel like a step up from GLM-5.3 [35].
- **Fireworks breach:** reporting says Notion was also affected, and that unlike Harvey, Notion's customers' data appears to have been exposed [36].
- **Vals CyberOffense:** Opus 5.5 averages 77.1% across three scenarios. In one run it completed every objective on a simulated corporate network with no human help, a first for the benchmark [37].
- **China data-center costs:** SemiAnalysis says a 100MW data center in China costs under $300M to build, excluding IT equipment. That is about a fifth of the US cost [38].
- **ChatGPT for businesses:** OpenAI launched Businesses in ChatGPT in beta. Businesses can manage their information there, and eligible ones can build customer-facing agents [39].

Sources

1. X article by @GoogleAIStudio
2. X post by @AminVahdat
3. X post by @ValsAI
4. X post by @_philschmid
5. X post by @perplexity_ai
6. X post by @cursor_ai
7. X post by @skalskip92
8. X post by @skalskip92
9. X post by @pushmeet
10. X post by @SemiAnalysis_
11. X post by @SemiAnalysis_
12. X post by @kimmonismus
13. X post by @CAIS
14. X post by @ValsAI
15. X post by @ValsAI
16. X post by @AnthropicAI
17. X post by @kimmonismus
18. X post by @kimmonismus
19. X post by @eli_lifland
20. X post by @tamaybes

21. X post by @TheRundownAI
22. X post by @ClaudeDevs
23. X post by @ClaudeDevs
24. X post by @GoogleDeepMind
25. X post by @GoogleDeepMind
26. X post by @GoogleDeepMind
27. X post by @DrJimFan
28. X post by @kimmonismus
29. X post by @MiniMax_AI
30. X post by @xeophon
31. X post by @GoodfireAI
32. X post by @GoodfireAI
33. X post by @Alibaba_Qwen
34. X post by @follistatindev
35. X post by @bclavie
36. X post by @steph_palazzolo
37. X post by @ValsAI
38. X post by @SemiAnalysis_
39. X post by @OpenAI