

TECHNOLOGY

Anthropic Accidentally Made the Perfect Commercial

How a 90-second ad encapsulates the AI industry

By Lila Shroff




Illustration by The Atlantic. Source: Claude and Anthropic / Youtube.

JULY 15, 2026

SHARE AS GIFT 

DISCUSS (123)

SAVE 

In an impressive challenge to the idea that no publicity is bad publicity, Anthropic is now promoting itself with images of disaster, death, and human suffering. The AI company's latest ad, which debuted during a World Cup match last week, is a perplexing 90-second artifact. The commercial opens with footage of a house engulfed in flames and the crackling sound of burning wood. From there, a series of dark images—rows of gravestones in a cemetery, a homeless couple sleeping on the street, laborers working in a West African  (123)—flashes across the screen. These stills are accompanied by a series of voice-overs. “Can AI be trusted?” “Who’s going to

hit the brakes if we need to?” “How do we really ensure that what we’re aiming to achieve really does benefit the majority of people?” The questions are left unanswered.

ADVERTISEMENT

Halfway through, the ad moves from dystopian film trailer to feel-good corporate ad. We see a photo of a mother sitting beside her daughter as a voice speculates that AI will allow her to be a better mom. Later, a medical worker uses a laptop covered in Anthropic stickers. “Maybe it’ll cure some great things,” a man says. A nature montage plays—there’s a humpback whale, grassy hills, and people hugging.

Finally, text flashes on the screen: “There’s hope in hard questions.” In the background, we see the ocean. And then Anthropic’s tagline: “Keep thinking.”

The commercial is tone-deaf. It’s one thing to mention the possibility of hitting the brakes on AI progress; it’s another to overlay the question on an image of dead Americans. The video is also just confusing. Does Anthropic think that AI is going to set the world on fire and leave us jobless, or that its chatbot, Claude, will lead us to a more human future where we spend our days watching whales by the sea? The company’s answer seems to be: *We don’t know*. The video has been mocked, obviously. Some people likened it to hypothetical ads by oil companies professing concerns about the environment. The head of one tech-media company posted, “I actually like the tagline but we don’t need the shots of *checks notes* burning houses and graves!” Meanwhile, OpenAI’s Sam Altman—who has a self-interest in undercutting Anthropic—chimed in to call the ad “satire.”

This is not the first time that a company has botched its advertising around AI: OpenAI has had its own share of embarrassing ads. By now, it seems like a rite of passage. Two years ago, Apple ran a commercial depicting musical instruments and art supplies being crushed by a hydraulic press. The ad was supposed to show off just how much creativity could be compacted into an iPad, but instead it seemed to symbolize generative AI’s bludgeoning of the arts. Apple subsequently apologized. During the Olympics that summer, Google ran an ad in which a young girl wants to write fan mail to the track star Sydney McLaughlin-Levrone—so her dad fires up Gemini and

asks the bot for help. The ad was highly unpopular, and Google pulled it from TV. If anything, Anthropic has had a comparatively good run until now. During this year's Super Bowl, the company aired a commercial dunking on OpenAI for putting ads in user conversations, which even Altman said made him laugh (despite his qualms with how the ad presented OpenAI).

Read: Watch Apple trash-compact human culture

For all the criticism, Anthropic's latest commercial captures the company's ambivalence about its own existence. Here is a leading AI company worth nearly \$1 trillion arguing that the very technology it is racing to commercialize may well imperil the world. "We founded Anthropic because we believe the impact of AI might be comparable to that of the industrial and scientific revolutions, but we aren't confident it will go well," the company wrote in 2023. In recent years, Anthropic has made a show of positioning itself as the ethical AI company. Earlier this year, it attracted the ire of the Trump administration when it pushed back against the Pentagon's demands for unrestricted access to its AI models over concerns that the technology could be misused.

The new advertisement is part of a broader campaign Anthropic is running to address the "hard questions" raised by AI's progress. In an email, an Anthropic spokesperson pointed me to a series of blog posts about the initiative: This spring, the company traveled across the country to better understand the public's views about AI. In the video's description on YouTube, Anthropic says that every question in the ad comes from a person that the company spoke with. (The spokesperson declined to comment further.)

The idea that Anthropic is worried about AI's negative impacts and wants to hear from the public is certainly better than the alternative. The trouble, though, is that as Anthropic has matured, the company's constant messaging about its anxiety over the future of AI has become less palatable and, or (123) nally, downright maddening. "Anthropic is at war with itself," my colleague Matteo Wong wrote in January. Since

then, the company's valuation has soared. As it develops more advanced AI models and races toward a public offering, the contradictions are becoming more difficult to square. However sincere Anthropic may be—Wong noted after several visits to the company that its employees can be “almost pathologically earnest”—it's no surprise that people are wary of the world's most valuable private company asking the “hard questions” about AI.

Read: Anthropic is at war with itself

The same tension is at play across the industry. OpenAI was founded as a nonprofit with a mission to “ensure that artificial general intelligence benefits all of humanity.” In the years since ChatGPT was first released, the company seemingly has prioritized financial gain instead. By now, it feels pointless for Anthropic and the rest of the industry to simply pose questions about AI ethics and safety. These questions are already consuming the public's imagination amid a mounting AI backlash. What matters is the answers. And on that front, AI companies have less to offer. Perhaps the commercial's tagline provides a helpful clue: For now, Anthropic will keep thinking.

[View Discussion](#) 123

ABOUT THE AUTHOR

Lila Shroff



Follow

Lila Shroff is a staff writer at *The Atlantic*.

Explore More Topics

Anthropic

About

Contact

Podcasts

Subscription

Follow

[Privacy Policy](#) [Your Privacy Choices](#) [Advertising Guidelines](#) [Terms & Conditions](#) [Terms of Sale](#) [Responsible Disclosure](#) [Site Map](#)

TheAtlantic.com © 2026 The Atlantic Monthly Group. All Rights Reserved.

This site is protected by reCAPTCHA and the Google Privacy Policy and Terms of Service apply