

Our tech overlords are planning for conscious AI to conquer the cosmos. What could go wrong?

Illustration: James Kerr / Scorpion Dagger

A new belief set is uniting some of the wealthiest men in the world
around a 'transhuman' future – actual humanity be damned

Eduardo Porter

Sam Altman, the chief executive of OpenAI, took to the Internet a few years ago **to propose** that homo sapiens would be the first species “to design our own descendants”. In his best case scenario, the “merge” between humans and artificial intelligence occurs at some point over the next 50 years. The alternative, where we remain simply human and the machines follow their own path, is more ominous. “If two different species both want the same thing and only one can have it – in this case, to be the dominant species on the planet and beyond – they are going to have conflict,” he wrote.

More recently, Elon Musk, the world’s richest man, who at one point last year was granted the power to reconfigure the US federal government, argued **on his social media platform, X**, that “it increasingly appears that humanity is a biological bootloader for digital superintelligence” – our role in the history of the cosmos reduced to that of

the low level code that boots up a computer before you can run sophisticated programs on it.

And Musk is on the tame side of the evolutionary proposition. According to [Silicon Valley lore](#), he once pushed back against Google co-founder Larry Page's claim that our next manifestation, to follow in the steps of the meat-and-bone humans you see walking about today, would necessarily have digital form in order to spread throughout the galaxy. (In fact, he recently testified in court that it was those concerns that prompted him to found OpenAI with Altman.) Meat and bones do not make for efficient interstellar travelers.

*Who cares if artificial intelligence obliterates
humdrum human labor when it offers us the
opportunity to transcend our body and conquer
the galaxy?*

It would be a mistake to understand these weird worldviews as an ultimately harmless take by techies who grew up on a diet of dystopian science fiction. The notion that we are approaching the end of the homo sapiens, as defined since Darwin's day, is coalescing into a durable body of belief among the elites at the helm of our technological future.

Their dreams are not all perfectly aligned. But like the folk stories and superstitions that have for ever revolved around more established religious traditions, the collection of far-fetched scenarios valley oligarchs are writing into our future exhibits the hallmarks of a religion in the making, a body of belief to confer a sense of cosmic transcendence and inevitability to their hi-tech project.

No death and an enhanced life: Is the future transhuman?

[Read more](#)

In their minds, they are on their way to build the next phase of humanity, a “transhuman” future. In this future, they can satisfy their desire for immortality and assert power over the cosmos as transhumans multiply and expand across the galaxy. Their ultimate goal: to execute on a techno-mystical dream to distill the essence of what it is to be human, consciousness and all, into bits of information to be downloaded as binary code on to some **non-biological substrate** such as a silicon chip, or beamed through space as electromagnetic waves.

The mythopoeic infrastructure assembled in and around San Francisco carries risk for humanity as we know it. It justifies steering technology along a path that is, at best, indifferent to the needs, hopes and aspirations of everyday humans in a quest to deliver a future that only looks like utopia to these masters of the universe. Who cares if artificial intelligence obliterates humdrum human labor when it offers us the opportunity to transcend our body and conquer the galaxy? The fantasy directs the technology: rather than building economically useful tools that can help humans expand their capabilities, the overlords of AI are sinking vast resources into a dream of building superhumans. These beliefs have pushed to the fore over the last quarter century, accompanying the advance of information technologies that have delivered enormous wealth and power to a new IT elite, one committed to science-based progress and hungry for transcendent meaning, but indifferent or even hostile to the propositions and moral constraints of organized religion.

“Silicon Valley has been a militantly secular space,” a prominent thinker about technology whose employer would be unhappy if he went on record told me. “It created a God-shaped hole, which it filled in its image.” Having rejected standard religious sources of purpose, they found an alternative path to provide their lives with significance via sci-fi transhuman dreams. Or as Musk observed in a singsong [post on X](#): “Atheism left an empty space. Secular religion took its place.”

While this newfangled cosmogony has been cobbled together at least since the early days of the Internet, it reached toward breathtaking new horizons on the shoulders of artificial intelligence, which opened up vast new possibilities for the transhuman dream. Douglas Rushkoff, a [critic of the technological oligarchy](#) and its ambitions, put it thus, referencing the 1980s-era satire featuring the first ever “computer-generated” TV host. “I guess AI makes the notion of having a [Max Headroom existence](#) plausible.”

Weird though the valley’s proposed utopia may appear, it fits a longer tradition of business titans with vast unrestrained wealth seeking to endow their endeavors with transcendent value. Henry Ford, as historian [Kati Curts](#) has written, also believed his calling was about more than transforming manufacturing to make cars; he believed he was on a mission to re-engineer the world to improve society.

Ford built Fordlândia, an attempt to create a harmonic social order supported by an industrial-scale rubber plantation in the Brazilian rainforest. Altman, Musk and the valley gang want to merge consciousness with AI and conquer the cosmos. The distance between these visions has mostly to do with the technological possibilities of their time. The proposition that they are engineering some utopian vision that humanity should be grateful for is not that dissimilar.

[View image in fullscreen](#)

Fordlândia in Aveiro, Brazil. Photograph: Joel Auerbach/Getty Images

As Nobel prize winning economist **Daron Acemoglu wrote**: “The handful of people unleashing this technology on the world are guided by an ideology of control (over humanity) and by a conviction that machines are uniformly better than humans.”

The danger, for the rest of us, is how the technological oligarchy’s aspirations will reshape the economies and societies of our present, as they redirect resources – capital, energy, minerals, water – to turbocharge AI and bring about the transhuman dream at the expense of healthcare, education or poverty reduction in the here and now.

While Americans are starting to show some signs of discomfort over the unrestrained appetites of this crop of AI moguls, the Trump administration has **shown few signs** so far of wanting to put in place regulatory guardrails and constrain their efforts in any way.

Future utopias on the menu

There are a variety of views in the valley about what a future humanity should look like.

Altman and Page are perhaps the most committed to the goal of merging humans with superintelligent technology and abandoning the flesh. Altman was an early subscriber to Nectome, **a valley startup** that proposes to retrieve information present in the brain’s anatomical layout and molecular details in order to replicate consciousness in the future. “I assume my brain will be uploaded to the cloud,” Altman **told** the MIT Technology Review.

Musk wants something a bit different, also spacebound but committed to flesh, enhanced by computers via something like his own brain-to-computer interface company Neuralink. Peter Thiel, of PayPal and Palantir fame, frowns on “just a computer program that simulates me”, but **is drawn to the techno-ideal** of “this radical transformation where your human, natural body gets transformed into an immortal body”.

And yet, the visions converge. Page, for instance, has suggested that rather than giving money to charity he might just give it to Musk. As he once **told Charlie Rose**, Musk wants to go to Mars to provide a backup planet for humanity to expand and that is a worthy goal to contribute to.

There are shared sources that provide some sense of moral purpose to the various flavors of sci-fi ambition. One of the core starting points is rather earthbound: the movement for effective altruism (EA), which seduced the technological elite with its appeal to unflinching rationality. Philanthropy, the EAs argued, was largely wasted by funding, say, the local library. Donors had to be purposeful, carefully directing their money to where it would do the most good for the most people. That is not an unreasonable proposition. It encouraged laudable efforts to, say, eradicate malaria in Africa, on the grounds that one could save a whole human life for a small fistfull of dollars. But it eventually departed from the needs of present earthlings.

[View image in fullscreen](#)

Peter Thiel after a meeting with Argentina's president, Javier Milei, at the Casa Rosada presidential palace in Buenos Aires, Argentina. Photograph: Matias Baglietto/Reuters

First, it was the longtermists, who emerged from effective altruism to argue that improving the world of the future was worthier than spending on the present. From there it took but **one small step** to move the goalposts to the cosmos: how about focusing on the wellbeing of myriad future transhumans populating the vast reaches of the galaxy in the far future? Maybe they will be of the flesh. Maybe not.

It's easy to get lost in the tangle of beliefs and aspirations – articulated and refined by academics like **William MacAskill** and **Nick Bostrom**, at university departments or thinktanks funded by the techno-oligarch's mushrooming wealth. They draw from unorthodox ethics, and from idiosyncratic readings of the laws of physics. The goal: to

justify the imperative to take humanity (or at least the most privileged part of it) where it has never gone before.

One of this crew's goals is to advance up the **Kardashev scale** – a measure of the amount of energy a civilization consumes – to harness the energy and acquire the technological capabilities needed to **transcend our biological confines**. Present day humanity, at the bottom of the ladder, doesn't even consume all of the energy of the Earth. Advanced civilizations, the thinking goes, are expected to consume all the energy of their star, at least, if not all that of the galaxy.

One of the earlier groups pushing for a transhuman future in the 1990s were the ultra-libertarian **Extropians**, which included leading intellectuals such as **Eliezer Yudkowsky**, Bostrom and economist **Robin Hanson**. Outlined in their core principles, they proposed "Boundless Expansion: Seeking more intelligence, wisdom, and effectiveness, an unlimited lifespan, and the removal of political, cultural, biological, and psychological limits to self-actualization and self-realization. Perpetually overcoming constraints on our progress and possibilities. Expanding into the universe and advancing without end."

Another, more recent branch, are the effective accelerationists. They have tried to conscript physics to their cause, arguing – controversially – that maximizing intelligent life is an imperative, because life is good at extracting available energy from the environment and dissipating it – increasing what is known in physics as "entropy".

As Beff Jezos – the online identity of **Guillaume Verdon**, one of the leading lights of the movement – **puts it**: "Effective accelerationism aims to follow the 'will of the universe': leaning into the thermodynamic bias towards futures with greater and smarter civilizations that are more effective at finding/extracting free energy from the universe and converting it to utility at grander and grander scales."

These aren't AI firms, they're defense contractors. We can't let them hide behind their models

Read more

In a philosophical twist that surely pleases Silicon Valley's billionaires, effective accelerationists argue for rampant techno-capitalism, unhindered by regulation, government and other nuisances, because this would maximize the consumption of the universe's resources, "**capture civilizational utility**", and dissipate the residue into the disorganized void.

The details of the dream don't actually make much of a difference. Because they all take us roughly to the same place. What matters now is whether the masters of the universe – invested in harnessing the energy of the stars, tempted by a moral calculus that posits that the wellbeing of the people of the present is of inferior value to the vastly more numerous humanoids of the future – will have the patience to care for the rest of us.

The signs are not great. Silicon Valley venture capitalist Marc Andreessen, for instance, wants to "ensure the techno-capital upward spiral continues forever". His **list of enemies** encompasses pretty much any person or idea that might stand against technological endeavor. That includes "sustainability", "social responsibility" and "tech ethics".

*If you create God but you own God you become
the dictator*

Jaron Lanier

Thiel is unusual in this crowd in that he is fiercely committed to an idiosyncratic variant of Christianity in which anybody standing in the way of technology, or governments that **try to tax him**, show up as the antichrist. But though he claims **little affinity with Andreessen**, he seems to have similar tastes. A diehard libertarian, he is contemptuous of government redistribution. **His philanthropy is about** for-profit investments in projects to further technological progress. Charity, as commonly understood, amounts to wasting resources that technologists will need to transcend our present. Musk has called empathy "the fundamental weakness of western civilization".

Regardless of the specific features of their transhuman dreams, the narrative crafted by Silicon Valley billionaires justifies their vast accumulation of power. As computer science pioneer and tech visionary Jaron Lanier told me: “If you create God but you own God you become the dictator.” And these dictators don’t seem to believe earthbound humans – most of us, at least – are particularly valuable. Questioned in February about the vast amounts of energy sucked up by AI, **Altman noted**, somewhat disparagingly, that “it also takes a lot of energy to train a human.”

[View image in fullscreen](#)

Jaron Lanier in his home in Santa Cruz, California. Photograph: Winni Wintermeyer/The Guardian

The flat-out indifference toward the rest of us is evident in **their frequent assessments** about what AI could **bring down upon us** – ending human work, building weapons of mass destruction, even bringing about **human extinction** in the service of making paperclips. **Palantir’s manifesto** notes that “one age of deterrence, the atomic age, is ending, and a new era of deterrence built on A.I. is set to begin.” Or as **Musk once put it**, before he changed his mind, launched xAI and merged it with SpaceX, “with AI we are summoning the demon.”

Yet they admittedly have no idea what they are doing. “People outside the field are often surprised and alarmed to learn that we do not understand how our own AI creations work,” Anthropic co-founder **Dario Amodei** wrote last year. “They are right to be concerned: this lack of understanding is essentially unprecedented in the history of technology.” Amodei has deep ties to effective altruism; his sister Daniela, Anthropic’s president, is **married to a founder of the movement**. Recently, though, they’ve both distanced themselves from it.

What’s particularly distressing is how unconstrained these moguls are, as they pursue the futuristic utopia they plan to build with their machines. Tech billionaires are plowing hundreds of millions into political campaigns, to fend off attempts at regulation and evade accountability lest their endeavors go awry. They want to make sure nobody butts in as they work to reshape

society. And they are largely succeeding – for now, no one with the power to stop them is butting in.

What is to be done?

How should society intervene? Does our political system provide the tools to help steer the process in a pro-social direction? Beyond the uncertain impact of technology on our future economic and social landscape, how should we address the narrow concentration of the fruits of these endeavors to build transhuman cyborgs with silicon brains?

The Trump administration has shown little interest so far in resisting the tech oligarch's fantasy. But that doesn't necessarily mean that the valley oligarchs' project of techno-domination is inevitable. Misgivings are emerging among the Maga base: The folks in rural Virginia who **push back** against datacenters hogging power and water supplies, **evangelicals wary** of a cosmopolitan elite claiming recourse to a tech-inflected higher authority.

Other signs of trouble are brewing for the AI project – from college graduates **booing commencement speakers** who extol AI, to Trump's brief moment of concern over the potential criminal capabilities of Anthropic's new Mythos model before **deciding not to regulate** the thing after all. In the latest **Times-Siena poll** from earlier in May, more than twice as many registered voters said AI is mostly bad, compared with those saying it was mostly good.

Perhaps the most forceful, pro-human position has come from the Holy Father himself. On Monday, Pope Leo published the encyclical Magnifica Humanitas, pushing back against the unfettered development of AI at the expense of jobs and social equity. "This creates a paradox of material progress and anthropological

regression that undermines the foundations of a just and stable social peace,” he added.

[View image in fullscreen](#)

Pope Leo attends the presentation of Magnifica Humanitas in Vatican City. Photograph: Alessia Giuliani/CPP/Shutterstock

One might also take comfort in the fact that the oligarchs’ dreamscape is so far-fetched. Ford and his civilizational dream again come to mind. Fordlândia today lies in ruins. A pointless water tower pokes into the sky from the banks of the Amazon, large decrepit houses in the American suburban aesthetic surround a lifeless playground and a long-empty swimming pool.

There are the ruins of nurseries, where as Federico Guzmán Rubio writes in his book *There is Such a Place*, Ford’s aversion to cows meant the children of workers were introduced to soy milk, shipped in from miles away. There are the ruins of schools where kids were taught about Henry Wadsworth Longfellow. What’s left is testament to the incongruous dreams of an oligarchy that overvalued its power and confused its appetites with the greater good.

AI companies will fail. We can salvage something from the wreckage

[Read more](#)

The AI-fueled cosmic fantasy is no less nuts. Forget the part where human consciousness is rendered in digital form, merged with AI and beamed across the galaxy. The ostensibly more down-to-earth proposition that **conscious AI** is not just possible but around the corner is in fundamental tension with our tenuous grasp of **what consciousness is**. Even more mundane objectives, such as getting artificial intelligence to train itself, keep getting pushed forward into the event horizon.

Perhaps this time too the outlandish claims will fade into irrelevance; the Star-Trek vision of people being dematerialized and **beamed up** and down around the galaxy will decay into some rustbound heap. Maybe the transhuman project will give way to a more or less recognizably human future with some cool new AI **plugins**. Maybe it can even be achieved in a way that serves our long-forgotten dream of equitable prosperity.

So far, though, our technological visionaries are pushing for something else, a future marked by vast concentrations of wealth and power, indifferent to the humdrum aspirations of the unwashed many. In the unlikely event that it succeeds in taking the essence of Page, Musk and their ilk aboard a silicon body to “where no man has gone before”, here’s hoping that they don’t destroy the world we know in the process.

- Eduardo Porter is a journalist focused on economics and politics. He is a Guardian US columnist and writes the newsletter **Being There** on Substack