

## Silicon Valley's 'great replacement theory'

*The Washington Post*, 12 Feb. 2025, by Eduardo Porter

5

It is no surprise that the appetite for oddball science fiction ideas that gave America everything from Scientology to Heaven's Gate would be irresistibly drawn to artificial intelligence. The image of Silicon Valley's enlightened CEOs hotly debating the merits of various paths toward a "transhuman" future — in which AI empowers humans to supersede our "biological substrate" and expand across the universe — is almost funny. But the vision loses much of its charm when you consider the power these men can wield. And it becomes downright terrifying when prominent members of this cabal are given the keys to the government of the most powerful nation in the world.

Take Elon Musk who is busy tinkering with the machinery of the federal government. Musk has a well-documented affection for "longtermism," the ethical foundation for the proposition that ensuring the future population of the cosmos by whatever-succeeds-present-day-humans is the moral imperative of the day.

The thought flows from a slightly tweaked utilitarian worldview: We should strive to maximize the well-being of the greatest number of people, no matter whether they live in the present or in the future. If you propose that we could reach the stars and beyond to populate the universe with bazillions of "people" living happy lives, you are left, morally, with no alternative but to devote yourself to ensure that this future comes about.

The Silicon Valley crowd is convinced AI has brought that prospect within clear view.

Our meat-based body is not the best vehicle to expand our consciousness beyond the galaxy. We die too quickly for interstellar colonization. We require lots of maintenance. AI can take our consciousness to the next level — maybe by living in a simulated reality in some piece of software. Or something. Musk loves this stuff. He has gone on record to say humanity is most likely already living in a simulation and has endorsed another of Bostrom's propositions: that humanity's overriding priority should be to mitigate the risk that some catastrophe prevents us from developing the techno-utopian version of ourselves that can populate the universe.

And Musk just made a \$97.4 billion bid to buy ChatGPT creator OpenAI.

He is not the only believer in our transhuman potential. Google's Larry Page has argued that digital life is the desirable next step in the cosmic evolution and that AI paranoia would delay the advent of utopia. OpenAI boss Sam Altman believes that "unless we destroy ourselves first, superhuman AI is going to happen, genetic enhancement is going to happen, and brain-machine interfaces are going to happen." Humanity, he says, "will be the first species ever to design our own descendants."

Musk at least appears to accept a case for government regulation to prevent AI from simply destroying everybody. A couple of years ago, the Future of Life Institute, to which Musk was a prominent donor, issued a letter signed by a bunch of intellectuals calling for a "pause" in big AI experiments lest we "risk loss of control of our civilization."

Another branch of posthuman enthusiasts, who call themselves "effective accelerationists," would rather let AI rip. Their argument is that the solution to killer AI is not regulation but rather a world populated by many AIs designed by competing firms. (537 words)

45