Ars on your lunch break: The fate we might be making for ourselves

Enlarge / Suck it, Skynet.

reader comments

17 with 14 posters participating

Share this story

Today we're presenting the second installment of my conversation with Naval Ravikant about existential risks. Naval is one of tech's most successful angel investors and the founder of multiple startups—including seed-stage investment platform AngelList. Part one of our conversation ran yesterday. If you missed it, click right here. Otherwise, you can press play on the embedded audio player or pull up the transcript—both of which are below.

This interview first appeared in March as two back-to-back episodes of the After On Podcast (which offers a 50-episode archive of unhurried conversations with world-class thinkers, founders, and scientists). As I mentioned in yesterday's article, my conversation with Naval led to a last-minute invite to give a related talk at April's TED conference. TED posted that talk to their site this morning, and if you feel like watching it, it's right here:

"How synthetic biology could wipe out humanity—and how we can stop it."
My talk focuses on the dangers that abuses of synthetic biology technology could lead to. Naval and I will tackle that

subject in our next two installments. Today, we focus on that time-honored Hollywood staple—super AI risk.

We both believe the greatest dangers from this lie many decades out. A big one might be the vast personal upside a super AI breakthrough could dangle in front of entrepreneurs and investors. The allure of private gains are wildly destabilizing when selfish actors can place bets that imperil the public good. Consider the 2008 financial crisis: for years, fat payoffs from toxic bets lined the pockets of corrupt financiers, who then faced no downside when things fell apart, and the rest of us picked up the $22 trillion bill.

The cost of an AI catastrophe would likewise be fully socialized. So would the costs of all-out nuclear war. But nuclear diplomats aren't in a position to privately rake in billions by dialing up systemic risks. If they were, we never would have made it through the Cold War.

But as Naval and I discuss, the bet-the-farm moments in AI's future could rest in the hands of entrepreneurs gunning for lavish IPOs. This would inevitably affect their risk tolerance. For related reasons, Naval is deeply skeptical of self-styled AI experts who dismiss even the faintest possibility that AI could ever pose an existential threat. Some take the extreme line that AI insiders alone have the credentials to assess AI risk. A parallel argument that only Goldman Sachs is clever enough to regulate Goldman Sachs is unlikely to get many takers.

If you enjoy this installment and just can't wait for parts three and four, you can binge them now by grabbing episode 45 from my podcast feed or my website. If you'd like to read a longer and broader article about existential risk (with a focus on synbio), I posted this to Ars earlier today. On a related note, my podcast's latest episode is one in which Kevin Rose turns the tables and interviews me. Our topic is

(what else?) existential risk. That same interview is running on Kevin’s podcast today.

This special edition of the Ars Technicast podcast can be accessed in the following places:

iTunes:
https://itunes.apple.com/us/podcast/the-ars-technicast/id522504024?mt=2 (Might take several hours after publication to appear.)

RSS:
http://arstechnica.libsyn.com/rss

Stitcher
http://www.stitcher.com/podcast/ars-technicast/the-ars-technicast

Libsyn:
http://directory.libsyn.com/shows/view/id/arstechnica

https://arstechnica.com/?p=1523275