Oxford Expert: 10% Chance Humans Extinct in a Century

I have not been blessed with a refined taste in cinema, with my favorite movie franchise being the Terminator series, especially the second and third, in which Arnie is in peak form. Alas, there's not enough space here to reminisce, so let's confine ourselves to the premise for the action.

On August 29, 1997, Skynet, an artificial intelligence system created by the US Defense Department, became self-conscious. Its programmers panicked and tried to deactivate it. Skynet defended itself by provoking a nuclear exchange in which three billion people died and the rest were enslaved or hunted down. Until John Connor organized the Resist…

Sorry, we must stop here as I've promised to talk about ethics.

Surprisingly, a minor academic industry exists whose goal is to solve the conundrums which might arise if (or when) Skynet or one of its buddies takes over the world. And this is just one of many apocalyptic scenarios which are on the table. The Global Priorities Project and the Future of Humanity Institute, both based at Oxford University, recently produced a [Global Catastrophic Risk 2016 report](#) which discusses some of the most likely ones.

It's less gripping than the Left Behind novels about the Second Coming (with titles like The Rapture: In the Twinkling of an Eye/Countdown to the Earth's Last Days), but, in its own dry, detached way, no less scary.

According to the Oxford experts' calculations, extinction of the whole human race is reasonably likely. Scientists have suggested that the risk is 0.1% per year, and perhaps as much as 0.2%. While this may not seem worthwhile worrying about,

these figures actually imply, says the report, that "an individual would be more than five times as likely to die in an extinction event than a car crash".

Tiny probabilities add up, so that the chance of extinction in the next century is 9.5% — which is worth worrying about. And of course, a mere global catastrophe, involving the death of a tenth of the population, is far more likely.

What sort of events do the futurists have in mind? The first of them has been on the front page of newspapers for several years: extreme climate change.

Then there is nuclear warfare, which would not only kill millions, but possibly trigger a nuclear winter. Pandemics like the Spanish Flu in 1918-19 have already killed millions. Natural events like the eruption of a supervolcano or a collision with an asteroid would be extremely challenging, as the dinosaurs discovered.

But what worries the futurists most is the risk of "emerging

technologies" such as Skynet in The Terminator. Oxford's Nick Bostrom, a philosopher from Sweden, is the leading light in the study of existential risk. In his recent book Superintelligence: Paths, Dangers, Strategies, he contends that artificial intelligence could become as powerful as the human mind, with a small, but hardly negligible, risk of something like Skynet developing. (Its message comes "highly recommended" by Bill Gates, which suggests that the world's richest man is not secretly planning to take over the world with a Microsoft version of Skynet.)

There are other runaway technologies which could destroy us. A killer microbe could be created with genetic engineering techniques which could wipe out whole populations. Colossal attempts to alter the climate with geoengineering techniques could backfire and turn the planet into desert or a snowfield. And then there are all the dangers which we foolishly don't fear because we don't even realize that they exist.

What, for instance, is the probability of Vogons showing up to build a hyperspatial express route through our star system? In A Hitchhiker's Guide to the Galaxy, it took slightly less than two minutes to demolish planet earth and only two people survived.

So here's where futurology stops and ethics begins: what should society do about massively destructive events with a low probability?

This is a question which is relatively recent, philosophically speaking. People began to pose it in the 1960s because of the threat of "mutually assured destruction" in a nuclear exchange, the imagined dangers of over-population, and climate change.

Nick Bostrom advises us not to wait for the worst to happen. He believes that "a moral case can be made that existential risk reduction is strictly more important than any other

global public good."

After doing a probability analysis of risk and future populations, he comes to the conclusion that "the expected value of reducing existential risk by a mere one billionth of one billionth of one percentage point is worth a hundred billion times as much as a billion human lives". This is difficult to comprehend, but the conclusion isn't: "the objective of reducing existential risks should be a dominant consideration whenever we act out of an impersonal concern for humankind as a whole".

In other words, we can never do enough to save humanity.

Personally, I find this blank check even scarier than supervolcanoes. It implies that governments should be empowered to tax to the max, spend freely, revise moral codes and restrict civil liberties to save humanity from invisible threats.

But is it sensible to entrust our future to statisticians? After all, calculations are only as good as the assumptions on which they are based. The old computing proverb, garbage in, garbage out, has yet to be disproved. It is easy to make enormous mistakes by moving a decimal point or neglecting to consider important inputs.

For instance, Paul Ehrlich confidently predicted that "hundreds of millions" would starve to death in the 1970s. This helped to create a world-wide panic over the "population bomb". To avert catastrophic risk, the Indian government embarked upon a campaign of compulsory sterilization which was an egregious violation of human rights and Western governments supported population control throughout the developing world.

But it never happened. Ehrlich and others had not anticipated the Green Revolution and falling birthrates.

And even at Oxford they make mistakes. Within days of issuing

the Global Catastrophic Risk 2016 report, the experts were eating humble pie. A mathematician reviewed its calculations and concluded that "the Future of Humanity Institute seems very confused re: the future of humanity". The authors had to correct their most startling statistic. It doesn't inspire a lot of confidence in the ethics of existential risk.

—

Image Credit:
Touchstone Pictures-Photo of Ben Affleck from Armageddon