




The Case for Tyranny

What is the strongest argument against libertarianism? Occasionally, someone asks that. Until recently, I didn't think there were any very strong ones.

But today I think there is at least one important argument against liberty that is hard to answer. This isn't exactly an argument "for tyranny" (as my title colorfully puts it — I have to attract clicks, you know) — there's no interesting argument for having the government send people to concentration camps, or prohibit people from criticizing the government, or prohibit private property, etc. But it's an argument for a pretty intrusive state. I'm going to explain it here to see if any readers can identify a good response to it.



I. You Should Care About Existential Risks

Assume that intelligent life is extremely valuable. If the human species continues for a long time, a large number of future people will exist, with valuable lives. Due to recent and expected future population growth, there will be many more future people than the people who have existed up till now.

However, there is a chance that something will extinguish the human species in the near future. Of course, something will *eventually* extinguish the human species — that's guaranteed. But it could happen either soon or in the distant future. If it happens soon, that will cut off all the enormous future value that the human species could otherwise realize. There is a good chance that then no more intelligent life would ever exist on Earth.

This means that actions or events that slightly increase the chances of the species being extinguished *soon* have very large expected costs. For example, if the species were to last for another million years, it's plausible that that would mean 100 trillion future lives.* Now, suppose some decision made by us today has a 1 in a million chance of extinguishing the human species in this century. That is the sort of risk that we would just ignore. If someone was worrying about a 1-in-a-million risk, we would laugh at that person.** But that risk would have an expected cost of *100 million lives*.

*Calculation: say the species lasts for 1 million years, with an average population of 8 billion (close to the current population; it's not clear whether future population will be higher or lower than this), and say that an average life is 80 years. Then we get: $(8,000,000,000 \text{ population}) (1,000,000 \text{ years}) / (80 \text{ years per life}) = 100,000,000,000,000$ lives. Obviously, these numbers are guesses. The point is just that the future has very large expected value, *if* the species doesn't destroy itself soon.



**Aside: this is how people react now to the threat of nuclear war. “Oh, don’t be silly, nobody would start a nuclear war. We’ll never have a leader crazy enough to want a nuclear war.” And “Don’t worry, a nuclear war probably wouldn’t kill *every-one*.” So we don’t care about the fact that the U.S. President has control of 4,000 nuclear weapons and the legal authority to launch a nuclear first strike at any time, for any reason. When I brought

this up once on Facebook, people laughed it off. We’re not close to a nuclear war now!

This is how our species is going to die. Not necessarily from nuclear war specifically, but from ignoring existential risks that don’t appear *imminent* at this moment. If we keep doing that, eventually, something is going to kill us – something that looked improbable in advance, but that, by the time it looks imminent, is too late to stop.

II. Freedom + Technology = Existential Risk

A. Freedom & crime

Until now in human history, political freedom has been an almost unmitigated good. The freer a society has been, the better off it has been in almost all ways.

But not *all* ways. Freer societies, for one thing, have more crime. The Soviet-bloc, communist countries had very low crime rates, for obvious reasons. Everyone was afraid of the state, which was constantly spying on the people. They didn’t have rights to privacy of the sort we are supposed to have in the U.S. So the state could effectively suppress crime. But this advantage was vastly outweighed by the general poverty and oppression of those societies.

But, that’s only because, up till now, individual criminals couldn’t do that much damage. Sometimes, a crazy person kills multiple people. In extreme cases (terrorists), maybe hundreds. In one extreme case, thousands. Until now, in order to kill more than that many people, you had to be a government.

B. The nature of technology

That’s going to change. I don’t know what future technologies we’re going to have. But there are three very broad features of technologies that we should expect to continue: More advanced technologies are (1) more *potent*, (2) more *efficient*, and (3) *pervertible*.

Meaning: As technology advances, it becomes possible for the person or group using it to produce larger effects, with smaller costs, where costs can be measured in money, in time, in effort and expertise. This makes it possible to do *more good*, if indeed one is trying to do good. But the power to do good is also the power to do evil, and the power to do *a lot of good* is also the power to do a lot of evil.

Example: When human beings invented jet airplanes, this enabled us to transport large groups of people a long distance, in a short time, with less expense than earlier transportation methods. To do this, one has to harness large amounts of energy – there will be a large amount of stored

chemical energy (fuel), which will be converted into kinetic energy during flight. But as an inevitable side effect, the technology can also cause enormous harm, if that energy is directed in a destructive direction (as in the 9/11 terrorist attack).

Virtually every technology that can do something great can also do something horrible, depending on the intentions of the person using it. The more potent it is for the one purpose, the more potent it will usually be for the other. If a knife can cut broccoli easily, it can also cut human flesh. If an AI system can be used to figure out how to save lives, the same AI can be used to figure out how to destroy lives. If genetic engineering enables us to create smarter, healthier, happier people, the same technology will enable us to create more deadly, difficult-to-kill diseases.

As I say, features (1) – (3) are true of almost every technological advance. Now project that into the future. Assume that technology will continue to advance for a very long way, and that the advances will continue to have these three features that virtually all advances so far have had. What's the logical extension?

The logical extension is that it becomes possible for a single person, cheaply and easily, to destroy everything.

Of course, I don't know *how* exactly, since I don't now possess the technology that people will have in the future. But that, logically, is where we have to end up, if technology continues to become more potent, more efficient, and remains pervertable.

As an example, perhaps we will develop efficient devices for genetic engineering. A very advanced device would be affordable by ordinary individuals, and would enable a non-expert to genetically modify an organism to have the properties desired by the user. This could be used, say, to create a maximally cute, long-lived, well-behaved kitten . . . or a virus with the transmissibility of Covid-19 and the deadliness of AIDS.



C. Don't trust people

Once that day arrives, we're done. If a technology is invented that enables any individual to destroy the species, the next day, the species will be destroyed.

This isn't a particularly pessimistic view about human nature — this prediction does not depend on thinking “humans are inherently evil”, or “humans are suicidal”. It depends on the well-known fact that humans are *diverse*. If there are 8 billion people, *someone* is going to have some crazy ideas. There's going to be *at least one* person who is going to somehow think that destroying everything would be a good idea. (See <https://fakenous.net/?p=1302>, <https://fakenous.net/?p=1333>.) There are assuredly some such people now — they just haven't killed all of us because they don't yet have the *ability* to do so.

III. The Tyranny Solution

That's the problem with freedom, in an advanced society. What can be done about it?

a. Targeted restrictions: The most natural thought is that we should tightly control *just* the really dangerous technologies, the ones that could be used to kill millions of people. So far, that's worked because there aren't that many such technologies (esp. nuclear weapons). It may not work in the future, though, when there are more such technologies. It may not be possible to anticipate all the ways in which someone might try to destroy the species, and so we may not know in advance all the things that need to be tightly controlled.

b. Defensive technologies: We'll build defenses against the main threats. E.g., we'll build defenses against nuclear weapons, we'll engineer ourselves to resist genetically engineered viruses, etc. Problem: same as above; we may not be able to anticipate all the threats in advance. Also, defense is generally a losing game. It's easier and cheaper to destroy things than to protect them. That's why we have the saying "the best defense is a good offense".

Example: during the cold war, the Soviets were considering building a missile defense system. (Something Reagan later wanted to do as well.) The U.S. Defense Secretary, Robert McNamara, told them: "That's fine. If you do that, we will not respond by building our own missile defense system. We will respond by building *more missiles*, to make sure that we could still overwhelm your defenses." The reason was that it was cheaper to build n missiles than to build a system that would destroy n missiles. The Soviets gave up on the missile defense idea.

c. Tyranny/the End of Privacy: Maybe in the future, everyone will need to be closely monitored at all times, so that, if someone starts trying to destroy the world, other people can immediately intervene. Sam Harris suggested this in a podcast somewhere. Note: obviously, this applies as well (especially!) to government officials.

d. A better alternative . . . ?

Someone please fill in (d) for me. Thanks.

PUBLISHED BY



Michael Huemer

Michael Huemer is a professor of philosophy at the University of Colorado. He is the author of more than seventy academic articles in epistemology, ethics, metaethics, metaphysics, and political philosophy, as well as six amazing books that you should [immediately buy](#).

[View all posts by Michael Huemer](#) →

📅 July 11, 2020 👤 Michael Huemer 📁 Politics

44 thoughts on “The Case for Tyranny”



Howard B

July 11, 2020 at 9:53 am

Whether you support Zionism or not, Israel is a security state that is also striving to preserve the freedoms and thriving at least of Jewish Israelis.

Their problems are limited to counter terrorism; however they are a peculiar kind of security state.

I imagine that small states where everyone is of a similar ethnic group are best situated to solve your problems optimally.

But that's just a guess and not exactly an argument, and is speculative

 **Tim J Reichert**

July 11, 2020 at 10:21 am

Good lord. The end of privacy is not tyranny. Why? Because humans have never actually had privacy. Consider the millions of years humans lived as hunter gatherers or if you want to just stick with Homo sapiens then consider the hundreds of thousands of years we spent as hunter gatherers. Zero privacy. Was that tyranny? No it was not.

At no time in my life (55 years) have I had privacy from the police. They have been able to look up my details, tap my phone records, know my location etc. Just watch Dateline ABC. These are crime stories and every one of them is solved by the police using their investigating powers to get phone records, follow people, compel witness testimony etc. The idea that humans require privacy is a fallacy and a myth created by bad people who want to hurt us. Privacy hurts us as a whole. It only helps bad actors. Radical transparency helps us as a whole. It does not harm good people. It does not require "tyranny." That's what the bad people want you to think. They want you to believe that if the government finds out things about you you're fucked. This is of course a bullshit attempt to get you to demand more privacy so they, the bad people, can continue to fuck you without the law stopping them.

We can and should vote for a system of radical transparency and we should all embrace it, otherwise, as our host has pointed out, we're toast. Hunter gatherers had no secrets. It is not in human nature to be secretive. We are a social animal that does best with radical transparency. We can actually have a government that knows everything about us but does no harm to us. Option D is radical transparency without tyranny. It also requires the end of capitalism and money but that's another story.

Pay no attention to me though. I'm just loon.

 **Anonymous**

July 11, 2020 at 3:44 pm

Do you have a source for "Freer societies, for one thing, have more crime"?

Im pretty sure the soviet block didn't really have low crime rates, even if the official stats indicate otherwise. And that's even without counting crimes of the state against its population.

 **William Peden**

July 12, 2020 at 5:30 am

There is some truth in this comment. It depends what we mean by "crime".

My parents travelled to the Soviet Union and even had their honeymoon in the satellite states (Hungary, Czechoslovakia, and Poland) and Yugoslavia. Violent crime was more or less non-existent, as was theft of other people's property. There was no graffiti.

On the other hand, they saw a HUGE amount of bribery, theft of state property, petty pilfering (especially of foreign currency) and so on. In the USSR, if you had a few copies of Playboy or some fancy nylon stockings, you could get young men/women respectively to do your bidding and break just about every regulation in the book. You could even get a half-decent meal in a Soviet restaurant!

 **Anonymous**

July 11, 2020 at 5:11 pm

It is possible that as the absolute damage that can be done by any single actor increases rapidly, the relative damage (to the size of human civilization) increases only moderately, stays constant or even decreases.

For example: the prehistoric man could only kill a few fellow tribesmen before being killed himself. Today's terrorist can kill a few hundreds or thousands, but the size of the human population is much bigger.

By the time a single actor will be able to exterminate all life on earth, human civilization might span multiple planets/galaxies.

 **Michael Huemer** 

July 11, 2020 at 6:13 pm

Yes, it would be very hard to extinguish an interplanetary civilization. Unfortunately, to a casual eye, it looks like we are much closer to having world-destroying technology (capable of killing all people on a planet) than we are to colonizing multiple planets. E.g., we have nuclear weapons, we have genetic engineering techniques that can modify viruses. We are very far from having permanent settlements on any other planet.

I also think that interstellar and intergalactic civilizations are very unlikely. If it was easy to develop one, then we would have already encountered some interstellar/intergalactic alien civilizations. Remember that many star systems in the universe are billions of years older than ours.

 **William Peden**

July 12, 2020 at 5:35 am

The issue is not "Is it easy to develop an interstellar civilization?" but "Is it easy to develop an interstellar civilization given a human or enhanced human level of intelligence?"

In other words, it's not the relative frequency of interstellar civilizations in general or among species that matters for assessing our chances, but the conditional probability for species with our level of intelligence or what we might attain via genetic/technological enhancements, including AI.

You could explain our lack of contact with other species' civilizations on the grounds that intelligence => human level is very rare. However, I don't know the relevant conditional probabilities, and I don't think that I can determine them by subjective feeling or objective rules (I'm not a Bayesian).

 **Tom Jenney**

August 20, 2020 at 2:15 pm

As far as the likelihood of alien civilizations and Fermi's paradox (if they're so likely, why haven't we found them yet?), a possible answer is the one suggested by sci-fi writer Liu

Cixin in The Dark Forest. Because technology can advance very quickly, the chances that an alien civilization will be a threat to your own are very high: even if its initial communications to the stars are broadcast with primitive radio technologies, there is a some chance that it will advance so quickly that they will be dominant. That dynamic might lead to two defensive strategies: 1) maintain strict interstellar radio silence, and 2) destroy any communicating civilizations before they can develop the technology to destroy you.

 **Anonymous**

July 11, 2020 at 5:36 pm

A few thoughts:

1. Historically, the State was THE best technology for a single actor, or a small group to bring ruin to whole ethnicities, cultures, classes etc. In other words, the State itself is a form of existential risk.

2. To the extent that avoiding existential risk is a public good (a.k.a. externality), the standard analysis of public choice applies. There are market failures, but there are also political market failures, and there's no a-priori reason to think a State will offer a superior solution to the existential risk problem. I think there's a strong case against a State on those grounds alone (see 1).

3. To the extent that avoiding existential risk is *not* a public good, a de-centralized market is expected to do a better job dealing with it. The market solution might result in reduced privacy, surveillance, etc. but it's guaranteed to be efficient (in the economic sense) and better reflect individual preferences and trade-offs.

It will also be an engine for innovation in existential risk management, something one would not expect from the State.

 **Michael Huemer** 

July 11, 2020 at 6:08 pm

#1-2: Certainly true, but that doesn't tell us how to address the existential risks from crazy individuals. We need a solution to all the existential risks!

#3: I think it's a public good. If I stop someone from destroying the world, that's a huge positive externality.

The point at the end sounded like my option (b), defensive technology. Again, I think defense is inherently disadvantaged compared to offense.

 **Anonymous**

July 11, 2020 at 6:59 pm

#1-2: "We need a solution to all the existential risks!"

Well, there might not be a perfect, or even a good solution, and trying to devise a solution which consists of a strong State may very well be worse than nothing (let alone worse than whatever a free market can come up with, externalities notwithstanding).

#3. To a large extent it probably is a public good, but it doesn't automatically mean a State will do a better job. And I think there's every reason to believe it will do a worse job.

Your option (b) might be the best we can get.



Michael Huemer 👤

July 11, 2020 at 6:10 pm

Btw, what's with all the "anonymous" comments? You all know that you can type in any name, right? You're not limited to either typing your legal name or just writing "anonymous".



Anonymous

July 11, 2020 at 6:47 pm

I didn't type anything, I guess it defaults to "anonymous".



J.D. Bertron

July 11, 2020 at 10:14 pm

You should come make this point at LOTR. I think you would learn a lot from Bob Murphy and Larken's views on this. There are so many wholes in this it's too easy to take down.

- Why assume that technology gets more powerful only for destructive forces ?
 - Why assume centralized monitoring would be more efficient ?
 - Why assume that society becomes more fragile ?
-



Michael Huemer 👤

July 13, 2020 at 4:30 pm

#1: I don't assume that.

#2: Can't answer this because I don't understand it. More efficient than what?

#3: I don't assume that. Not sure why assumed I assumed that.



Anonymous

July 12, 2020 at 2:03 am

I think this is sound but is it really surprising? The best and most natural case for liberalism is deontological and when the long-term is on the table, so to speak, any deontological constraints are going to be dwarfed by the consequences, unless one accepts that it does not improve the world to ensure that more well off people are brought into existence than not, other things equal.

 **Peter**

July 12, 2020 at 2:21 am

A very similar argument is made by Persson/Savulescu in 'Unfit for the Future'. Besides massive surveillance, they advocate human moral enhancement to reduce the probability of people doing evil stuff. If it should be possible in the future to (e.g.) genetically alter humans in such a way that they do not even think of destroying humanity, we may not need a surveillance state. Sounds very speculative to me, though.

 **William Peden**

July 12, 2020 at 5:44 am

In principle, you could also do it via enculturation rather than genetic means. We already do this with children: we train them to be disciplined and socially responsible.

Interestingly, it's also people's default response to the problem of bad AI: "Just program the AI so that it's not evil". Obviously, exactly how to do this is difficult in the AI case, let alone humans. However, it's conceivable that it would reduce and threaten liberty less than e.g. everyone having a chip in their brain which put them in a coma as soon as they were about to misuse the easily available ER tech, because the latter could be used for all sorts of other things e.g. causing a coma if people pursued revolutionary actions.

 **Michael Huemer** 

July 13, 2020 at 4:33 pm

Evolution genetically programmed us for survival and reproduction. There are still suicides and voluntary celibates.

Culture also trains us to follow laws, respect people's rights, and not murder people. There are still murders.

Those things explain why suicide and murder are *rare*. But they can't make it rare *enough* for us to survive in the long term, if 8 billion people can each destroy us.

 **Matthew Skene**

July 12, 2020 at 7:55 am

It would probably take a lot of different measures to cut down on the risks. One of those might be heavy liability for producers and distributors of materials. If gun sellers and manufacturers were liable for how people used their products, we would have the greatest background check system ever, created by the suppliers, in place overnight. They would insist on knowing everything about a buyer to ensure they wouldn't be on the hook for wrongful death. Similar consequences for distributors of new technologies would result in efforts by providers of those technologies to create methods to help ensure crazy people didn't get their hands on them.

Moral progress would also help here. Improving how people treat one another would decrease the number of crazy people who might be willing to do this. This seems to be a race.

Can we keep people good enough and responsible enough with dangerous things long enough to spread out far enough that the remaining dangers don't kill us?

The idea that technological development outstripping moral development means a species is doomed over the long run does provide a good answer to Fermi's paradox. Either that is impossible, or else part of moral development discourages either spreading out or interacting with species like ours.

 **A Country Farmer**

July 12, 2020 at 3:10 pm

This is an important topic. A nihilistic and derange Unabomber with the ability to create deadly viruses is a scary thought.

One counterargument is that it's predictable that the government's defense against such an attack will be inept; whereas, the free market might more quickly find a solution (e.g. vaccine, etc.). By analogy, it's better for everyone to have the possibility to own a gun rather than the government owning all guns because the government can't react to crime quickly enough. If everyone could own a DNA manipulator and there's a fertile market for bioengineering, it's more likely we can stop such existential threats.

 **Nathan N**

July 12, 2020 at 8:22 pm

d. Get MUCH better at moral education (or otherwise change people's psychology to be averse to causing existential catastrophe).

 **Yujiri**

July 12, 2020 at 8:29 pm

I think A Country Farmer and a couple of the "Anonymous" people have made excellent points (especially the one with 3 numbered points), but there are a few important ones I don't think have been mentioned.

Time-scaling: Just because we aren't strict deontologists doesn't mean we value a future consequence *equally* as an imminent one. I think there's good reason to say consequences closer to the present have more moral weight, basically because of situations like this: if we don't believe in time-scaling, it would seem every effort ought to be directed toward accruing power for benevolent actors rather than doing immediately beneficent things.

More importantly: moral conflict. That is, arguing that it would be right to install *yourself* as a tyrant to keep the human race alive != arguing that one should allow *others* that power. We pretty much always think our own intentions are good, but that doesn't mean *having a state* (not run by the person doing this thinking) is any better an idea than it was before these considerations. (As the others pointed out, a state is *far* more dangerous as far as existential risks than non-state actors.)

We'd also have to consider that from the perspective of one of the people you (the benevolently rights-violating dictator) rule over, there is little reason to see you any differently from any other tyrant. If your people were to rebel, you would not even be able to raise self-de-

fense arguments; you would have to also defend whatever you did to put down an uprising by drawing solely from the riskiness of human freedom.

Finally, I think you might be forgetting to factor in the value of the freedom of the future lives. If human freedom is expected to reduce the number of future human lives by 100 trillion, but the tyranny required to prevent that risk reduces the average quality of life by 50%, it's costing not just 50% of the value of existing human lives, but also 50 trillion future lives.



Michael Huemer 🧑

July 15, 2020 at 5:35 pm

The government creates existential risks. But it does not create *more* risk than individuals. If a technology existed by which any individual could destroy the human species, I believe with near-100% probability that the human species would be gone within a year. If that same technology belonged only to the largest governments in the world, that would not happen. The expected life of the human species would go up by many orders of magnitude.

Almost everyone wants the species (and themselves) to survive. There is just a tiny portion of humans who would not want that. That's why, if the decision as to whether to destroy the species is given to every individual, we're going to be immediately dead — but if the decision is given to a large group of people — any large group (as long as it's not selected for craziness, like the group of "all the world's terrorists") — then we're safe for a long time.

Anyway, I don't think the question is whether there should be a government, since there already is one, and it is unlikely to go away soon. The realistic possibilities, at least for now, involve different policies that the government could pursue.



Yujiri

July 19, 2020 at 5:58 pm

I see your point. By collectivizing decision-making power, states are less prone to extreme aberrations than individuals given the same power, even if states are much more malevolent for the same number of actors.

It's a pretty good concern. I don't think it's enough to make a good argument for allowing or attempting tyranny because of the other points, though.

I could see a future, if individual power became a really big threat, where a free society might find ways to give up privacy without giving up individual agency? I do libertarians tend to underrate that distinction.



Pseudonymous

July 12, 2020 at 11:33 pm

This seems like a false dichotomy. Why can't we have a liberal-but-for state, which has strong government restrictions in the relevant areas for preventing the spread of these pervertible technologies but doesn't interfere in other totalitarian ways? Perhaps it would collapse into a full-blown totalitarian state.



Eric Hanneken

July 13, 2020 at 9:22 am

This post reminds me of Bill Joy's 2000 essay, "Why the Future Doesn't Need Us" (<https://www.wired.com/2000/04/joy-2/>).

As you say, humanity may be fated to destroy itself. The best counterargument to the tyranny solution is that we shouldn't expect it to delay the end of the world for long, if at all. There may be no action that will save a risk-adjusted 100 million lives.



Yujiri

July 13, 2020 at 5:29 pm

I posted a lengthy reply here that got accepted but didn't show up. At first I thought it might just be taking a minute but it's been a day, and trying to post it again tells me I'm trying to post a duplicate comment. It isn't in the RSS feed either. Halp?



Michael Huemer 

July 15, 2020 at 5:24 pm

Approved! Was incorrectly marked as spam.



Anonymous

July 14, 2020 at 1:59 pm

Just wondering—why do the people of this future government need monitoring instead of the destructive technology itself?

Speculation: most (all?) future super-destructive technology will involve an essential connection to digital/computerized stuff. This connection can either be through the digital/computerized process of creating the technology or as a digital/computerized component in the technology.

It seems possible that advancements would allow something like a "cookie" to be installed on every digital/computerized thing that has a connection to a device with super-destructive capacity. These cookies would constantly monitor these technologies. If they are put to bad use, maybe an alert would be sent out, the machine would break down, etc. In short, something would be done.

The focus, then, is on monitoring technology, not people. Of course, "potentially destructive" would have to be defined in an excessively over-broad way—same with "bad use"—to get around the problems you mention in III.a.&b.

Admittedly, this would lead to a kind of tyranny. Many, if not most or all, devices would have cookies constantly monitoring their activities. But the possibility of self-monitoring and self-reporting technology might be preferable to complete, good old-fashioned tyranny. If the cookies are designed well enough, they may even be inaccessible to those in power with interests in monitoring peoples' activities.

Of course, I have no idea whether something like this is possible, but it doesn't seem to be completely crazy and appears to fall within the realm of imaginative possibility presupposed by the post.



Michael Huemer 👤

July 15, 2020 at 5:27 pm

Thanks, good suggestion. Unfortunately, I think people would figure out how to disable the security software. Or they would disconnect the computer from the other, destructive part of the technology.



Brian

July 16, 2020 at 8:05 am

Two things.

1.) Doesn't the arms-race argument here contradict your argument in your book that militaries would shrink out of existence under anarchism? In your book defense is less costly than offense, but here you're saying offense is less costly. You must have some argument for why an arms-race wouldn't happen or would be very unlikely under anarchism, so wouldn't that also apply here?

2.) I think the premise that technology creates or will create existential threats overnight is weak. I agree with the 3 features you listed, but it seems like we lack evidence that a given technology, while it's being developed, goes from being low-threat to existential threat overnight.

It seems like we would acquire evidence of this sort of thing with a much lesser global offense (and presumably much easier to accomplish) than wiping out humanity – like the technology to wipe out some major global infrastructure like the internet or financial markets – before an existential threat, and actually have some evidence that such rapid development could happen, be accessible to the masses (due to high development costs), and have such a massive impact, before considering conceding anarchism.



Vincent Flink

July 16, 2020 at 9:28 am

Counterargument: although best single alternative could very well be “pretty intrusive”, most of ways in which state can be intrusive would have bad consequence even for x-risk mitigation. Some might themselves constitute x-risk. Since epistemological situation bad – no intrude!



Jan

July 16, 2020 at 5:18 pm

It's probably even worse than you think. See Bostrom's paper on exacty this topic:
<https://www.nickbostrom.com/papers/vulnerable.pdf>

 **Thomas**

July 16, 2020 at 7:26 pm

If a hypothetical technology can create extremely virulent viruses, should we not assume it could also create extremely successful cures to said viruses too?

Nonetheless, does the example (genetic modification) you posit demonstrate an asymmetry pertaining to goods and harms a technology creates or rather a proportionality between the two? Because if a single person could create irreversible pandemics that would vastly outweigh any good it might procure. But if, using the same technology, the pandemic could be easily reversed, then the goods would now outweigh the harms. Therefore, there's disproportionality here.

 **Ryan Carey**

July 17, 2020 at 5:51 am

Is there a reason you don't cite <https://nickbostrom.com/papers/vulnerable.pdf> ?

 **Jon**

July 23, 2020 at 3:08 am

The solution would be to support a state structure that will raise humanity to not need surveillance. If you raise 'em right, you don't have to worry about 'em. Of course the difficulty is fettering society only just enough to stumble enough around to have cultivated a satisfactory range for rich creativity and genuine love but to also find its way to not needing a state. It's a perilous and scary ride for sure. It seems very unlikely that will happen without God or something like it.

 **Mand**

July 24, 2020 at 7:21 am

Hey, Huemer! Not an ancap here (rather an anarchist without adjectives, from what you might call left-libertarianism), but my solution did begin from capitalism.

You see, on our current world, surveillance is already pretty high, although it's not used with such good intentions. Facebook, Google and other social networks manage to get access to a lot of our information with ease, even understanding our personalities and manipulating our emotions (I can source all of these statements if you'd like me to).

This would allow for, within ancapistan, more popular digital service companies to gather data on the minds of dangerous individuals, perhaps even before they know they're dangerous, rate them based on danger and increase surveillance on those most perilous to the interests of those companies. These companies may even manipulate individuals as to reduce their risk factors, thus reducing the probability of this happening at all.

This would allow them indirect control over whether people do this, and much like governments they are less likely to support "extreme aberrations" as one commenter has called

them. This solution of mine does, however, still significantly reduce human quality of life, but it is more desirable than a state. There is still a better solution, though: it's to manage those same companies democratically, ensuring that data will be used more responsibly and putting further barriers against the activation of the threat. This is because a co-op has more barriers against a leader deciding to use such a dangerous technology than does a hierarchical model.

In that sense, mutualism would be an imperfect, but already great solution to the problem you've pointed out. It has the surveillance suggested, without being nearly as coercitive, is more invisible and the data gathered is less likely to be used against the interests of people.

A friend of mine also suggested using EVEN MORE information technology in order to make this surveillance pretty much identityless to those performing it. A.I and/or algorithms and such, to much like facial recognition be able to recognize an action regardless of who is performing it, in a sense allowing people to be surveilled in an automated manner without anyone being able to harm them or even know who they are in any way before a dangerous action is planned or in course, perhaps even making that action impossible upon detection.



Anonymous

August 6, 2020 at 6:16 pm

As environmentalists are so fond of pointing out, "there is no Planet B." If humanity wants to stay around a long time, we will need to fix that—space exploration, with a view towards colonization, must become a top priority of human civilization.



Other Anon

August 15, 2020 at 3:01 pm

I'll quote from a fictional character (a time traveller who has seen human extinction in the future and explains why he doesn't try to prevent it; at first he cites technical reasons but then he ends with this quote):

"To be completely frank, I don't care much for humanity anyway. If you look at any part of its history, it's always nasty and terrible. An endless cycle of violence, oppression and misery. I fuck young girls and keep you as slaves, but I always offer you a painless way out. Most of humanity was never this noble. They never had a problem torturing people who wished they were dead. Why would I want to save a species like that? All it would do is make me responsible for their future victims. Not to mention, I wouldn't gain anything in return for my effort."




Anonymous

August 16, 2020 at 7:53 pm

I want to make sure humanity survives for the same reason a mother bear wants to defend her cubs. I'm a product of natural selection and have no desire to oppose my deeply ingrained instincts.

Other Anon

 August 17, 2020 at 1:13 am

We are all products of natural selection with deeply ingrained instincts. For instance, I have the deeply ingrained instinct to oppose those who want to force tyranny on me. Yet this is rationalized in the name of species-wide collectivism. Unfortunately even Huemer has now been infected by this meme.

I don't see a good reason why I should project my empathetic instincts on all of humanity, instead of, say, their future torture victims who will endure fates worse than death and not be allowed to die.

Humanity includes all your enemies. Everyone who insults you daily on the internet, everyone who wants to destroy your rights and liberties, everyone who wants you to be tortured in Hell for ever and ever. They're all humanity. I see no reason to treat them as though they were my cubs. They're clearly not my cubs.

But your mileage may vary. Feel free to terraform Mars – on your own dime and without forcing tyranny on me or the people I care about.

 **Anonymous**

August 17, 2020 at 8:21 am

I accidentally posted may reply to you as a top-level comment; you can find it below.

 **Anonymous**

August 17, 2020 at 8:08 am

Instincts can't be justified logically, because reason and science tell us what is, not what we should do—morals and ethics are human constructs. In any case, I agree that humanity includes a sizeable contingent of SOBs, but I can't help thinking that had I been born in slightly different circumstances I might have been one. By the standards of some I probably *am* one, and so are you, and so is everyone else. So I identify with them.

(If you want a longer defense of humanity along similar lines, I recommend the ending of Heinlein's novel "Have Space Suit—Will Travel").