Analysis

What worms can teach us about the AI apocalypse

It's good to treat children for gut parasites, say 'effective altruists' — but be wary of killer robots too

BY Tom Chivers

Tom Chivers is a science writer. His first book, *The AI Does Not Hate You* is out now.

August 19, 2020

August 19, 2020

Filed under:

Artificial intelligenceDiseaseEffective altruismepidemiologyMedicineRisk

Share:

Children line up as a teacher gives them deworming tablets. Credit: NOAH SEELAM/AFP via Getty Images

If you're a parent, you're probably familiar with threadworms. If you're not a parent, then you may not have heard of them at all. They're one of those secret joys of parenthood that no one mentions until they turn up — grim, pale little creatures, wriggling in a child's poo — whereupon you learn that almost every other family has gone through the same thing. Your parents may well have gone through it, with you, and then just … never spoken of it again, because eeerrgh. I almost feel like I'm breaking some parenting *omertà* just by mentioning them. But I need to, because threadworms can help you understand why we need to worry about artificial intelligence.

They're unsettling, disgusting things, threadworms: or they were to me, when I first came across them. But – luckily – they're easily treated. A cheap, over-the-counter medicine, mebendazole, kills the lot. You have to dose the whole family, clean your linen, and dose again two weeks later, because the eggs (god, I'm shuddering just typing this) survive for a couple of weeks; but basically there is a simple, nuke-them-all-from-orbit cure.

Threadworms are just one species of intestinal worms, also known as helminths; there are hookworms, pinworms, tapeworms, whipworms, others. And they're common — not just in north London, but around the world. The US Centres for Disease Control estimate that *Ascaris lumbricoides* is present in the intestinal tracts of about 1 billion people around the world. Usually, the infection has no symptoms, but in children it can lead to malnutrition, slow growth and impaired learning.

Luckily, again, these helminths can be killed with the same drug, mebendazole, or a related one, albendazole. Just as with the threadworms that cause itchy bottoms in the rich west, the *Ascaris* that cause childhood stunting in the developing world can be nuked, effectively and cheaply. A course of albendazole costs a few cents.

More from this author

Coronavirus is the wake-up call we need

By Tom Chivers

Dosing someone who's been diagnosed with helminths, then, is a no-brainer. It's cheap and easy and effective. But over the past 20 years or so, there's been a row going on about doing more than that: about giving the drugs to entire schools' worth of children, regardless of whether they've been diagnosed with worms or not.

In the late 1990s, two economists, Edward Miguel of Berkeley and Michael Kremer of Harvard, ran a study. They took 75 schools in Kenya, and gave albendazole to all the kids in 50 of them. Their results were impressive: absenteeism dropped by a quarter; children did better at school. The benefits also seemed to spread to nearby schools which hadn't been treated, presumably by disrupting the spread of worms in the region. The whole thing cost $3.50 per pupil.

Where it got *really* interesting, though, was a follow-up study, a decade later. It found that the children in the schools not only did better at school — when they left school, the benefits continued. They earned more — about 15% more, on average — than those who had not been dosed. That is a huge benefit for a tiny cost. So everyone — including the WHO — got on board with deworming programmes. (Apparently, Cherie Blair even dressed up as a giant worm at Davos and chased people around.)

Then it all went a bit crazy. Two epidemiologists at the London School for Hygiene and Tropical Medicine reanalysed Michael and Kremer's data, and found errors in the code, and some missing information, which led them to severely downgrade the findings: notably, the effect on surrounding schools disappeared. (It's worth noting that there's no suggestion of malpractice, and the LSHTM team and others have been full of praise with M&K for sharing their code and data so fully.) They also noted that the study wasn't properly randomised or blinded. Then the well-respected Cochrane Collaboration released a report saying that there was little evidence that deworming programmes show benefit. Now a new study, again by M&K, claims that the economic benefits are still continuing.

suggested reading
Are humans heading for the knacker's yard?
By Peter Franklin

If you want to read about the whole saga, which came to be known as the Worm Wars, I recommend Tim Harford's thoughtful take on it, and Ben Goldacre's more sceptical one. There is a lot of back-and-forth, with defenders of M&K saying that the statistical techniques used by the LSHTM team made it almost impossible for the programmes to find a benefit. "To be quite frank, you have to throw so much crazy shit at Miguel-Kremer to make the result go away that I believe the result even more than when I started," said one professor of public policy. On the other hand, one researcher into evidence-based charity told me "no one believes this but the economists".

This all came up when I was writing my last book — which, somewhat incongruously, was about artificial intelligence, and whether it would destroy the human race. The two topics don't necessarily seem connected, but they are.

The effective altruism community is a loosely affiliated group of people and organisations who try to find the best way of donating and spending charitable money, so it will do the most good in the world. For instance, imagine you want to help people who have visual impairment. You're deciding whether to donate to a guide dog charity, or to a charity which provides surgery for trachoma in the developing world. As the effective altruism charity Giving What We Can points out, it costs about $50,000 to breed and train a guide dog, but you can restore many trachoma sufferer's vision with an operation costing about $1,000. "For the same amount of

money as training one guide dog to help one person," they say, "we could instead cure about 50 cases of severe vision impairment." That seems a fairly straightforward decision.

suggested reading
The inhuman politics of artificial intelligence
By jon cruddas

Sometimes it's not so clear. Is extending someone's life by a year worth more or less than providing secondary education for a child? Is funding research into a Covid vaccine more or less effective than supporting democratic movements in authoritarian countries? But often, since most of us give to charity in fairly un-thought-through ways, it's easy to improve the good we do. The simplest method is what William MacAskill, a professor of ethics at Oxford and founding member of the effective altruism movement, calls the "100x multiplier": if you donate to a charity that works in the developing world, you can roughly expect to do 100 times as much good as if you donate to one that does a similar thing in the rich West.

Effective altruist organisations such as GiveWell and Giving What We Can give lists of the most efficient charities to give to: the standout ones are about malaria. GiveWell estimates that roughly every $2,300 spent on bednets coated in insecticide will save one child's life.

What I find interesting, though, is that both organisations recommend the Deworm the World Initiative and other deworming programmes. The effective altruism movement isn't unaware of the Worm Wars controversy: GiveWell discusses it all in two lengthy, stats-heavy blog posts. They mention various reasons why the Cochrane review might have missed something — notably, it only looked at weight gain and school performance, not future earnings. But they also note that the children Miguel & Kremer's research looked at had abnormally high levels of infestation, and that there are no obvious short-term mechanisms which could explain the long-term benefits — which should make us less confident that there's anything real going on.

But they still recommend it — because, in short, they think *the gamble is worth it*.

There's a concept in probability theory called "expected value". That is, how much, on average, you can expect some action to return. Say someone offers to flip a coin with you: on heads they pay you £2, on tails you pay them £1. Your expected value is (2-1)/2, or 50p: if you played the game lots of times, you'd expect to make an average of 50p per throw.

suggested reading
Can poorer nations survive Covid-19?
By Daniel knowles

With deworming, effective altruists are doing exactly the same thing. There's a strong possibility that deworming has no effect whatsoever on future earnings, or at least a hugely reduced one. But there is a small possibility that it has a huge, overwhelming impact, at a very low cost. The expected value of deworming — the impact multiplied by the probability, minus the cost — is huge, even taking into account the possibility that it doesn't work at all. GiveWell assumes that its impact is 100 times less than the Miguel-Kremer study claims, and, even at that level, finds that it is cost-effective.

This is all very straightforward and I think most people would get on board with it. But the same reasoning leads you to stranger places, which is where AI comes in.

Nick Bostrom, the philosopher and AI researcher, points out in his book *Superintelligence: Paths, Dangers, Strategies* that there are a lot of humans who are yet to live. I discussed exactly how he works it out in an earlier post, so I won't go over it again, but he ends up with 10^58 humans (one followed by 58 zeroes; a big number) before we run out of universe. That would be a very great positive — assuming that we think that human life is a positive thing, and that we would, all else being equal, prefer more happy lives than fewer.

So ensuring that humans don't go extinct (a rogue AI or a bioengineered pandemic are considered the most likely ways that will happen) has a *huge* potential upside: quintillions and septillions of lives, unthinkably huge numbers. So even if there's only a one-in-a-billion chance that he's right, or one in 10 billion, it doesn't matter. Those numbers are so huge that the expected value of anything that might improve the chance even by a fraction of a percent make the expected value — the benefit times the probability — spectacularly large.

So, for large parts of the effective altruism community, the most effective way to do good in the world is to donate to organisations that will reduce the chance of an AI apocalypse, and to ensure that any superintelligent AIs we *do* build will work to help humanity survive, in a form that we would want.

More from this author
Are economists really this stupid?
By Tom Chivers

This sort of reasoning can take you to some very strange places, but it'd be too easy to wash your hands of it and say it's all crazy. It's not; if anything I rather admire the tendency to follow the reasoning where it leads, and not simply pull up short and say "OK now we're getting weird." My own instincts are that concern about existential risk is pretty sensible and that AI is a reasonably likely way that we could blow ourselves up. And most effective altruism is of the much more worldly, small-scale "antimalarial bed nets" calibre.

This latest Miguel/Kremer study, for what it's worth, doesn't seem to have shifted opinions much in the Worm Wars. The evidence-based charity researcher I spoke to points out that it is essentially a continuation of the previous studies, so all the problems that existed in the first studies are still there now. More than that, it follows the same group of people as the previous ones — it would be frankly weird if they had all suddenly become poor, after years of being richer. "Same shit, different year," the researcher said, succinctly. More importantly, this person says, there is a chance — which as far as I know has not been studied — that mass dosing with deworming drugs could cause worms to evolve a resistance to those drugs, like antibiotics create antibiotic-resistant bacteria. It is not inconceivable that the actual impact is negative, which would affect the calculations somewhat.

Miguel Kremer, meanwhile, has won a Nobel prize in economics. "So we know who won the Worm Wars," my source says. "I'm sure *that* will end all the arguments;" that last said with heavy sarcasm.

I've broken the *omertà*; if you don't hear from me again, I've been silenced by the parent mafia. Only parents and developmental economists are allowed to talk about de-worming. But now you know how itchy bottoms relate to the AI apocalypse.

Join the discussion

Richard Pinch August 21, 2020

I think the ball is in his court. Read more

M

Mark Corby August 21, 2020

Nothing Incidentally how did your discourse with Chris Martin end? I had you at 40-15, but must have missed the last game. Read more

Richard Pinch August 20, 2020

Slightly more seriously, though, one could imagine a Net Present Value with a discount rate of say 3%. So each future generation is "worth" half of the preceding one, and the total NPV of a human population of one billion for ever is still only equal to about two billion current-day lives. Read more

To get involved in the discussion and stay up to date, join UnHerd.

It's simple, quick and free.

Sign me up

Tagged Artificial intelligence, Disease, Effective altruism, epidemiology, Medicine, Risk

Explore

- Groupthink
- Capitalism
- Flyover country
- Faith & Meaning
- Confessions Podcast

Quick Links

- About UnHerd
- Columnists & Writers
- Subscribe for free

Contribute to UnHerd

We welcome applications to contribute to UnHerd – please fill out the form below including examples of your previously published work.
Please click here to submit your pitch.