

SAM HARRIS INTERVIEW PART FOUR

NOTE: This transcript was generated by the service Rev.com, and though it's quite good, it is imperfect. If you would like to quote anything from this episode you are welcome to! But please find the point in the audio where the quote originates, and verify that the transcript is accurate. If Rev's transcript is inaccurate, we ask that you only post or publish a version of the quote that you deem correct (because your ears are awesome!).

Hello again, Ars Technica listeners. This is the fourth and final installment of a four-part interview with neuroscientist, New York Times bestselling author, podcaster, and controversial public intellectual, Sam Harris. We'll pick up with the uplifting theme we left off on yesterday. Which is to say, high tech weaponry that tomorrow's nihilistic, suicidal mass murderers might just use to kill very large numbers of us.

Before we get started though, a quick note of orientation. As those of you who listened to last week's interview with UCSF neuroscientist Adam Gazzaley know, I originally thought this podcast series would be a limited set of just eight episodes connected to my latest science fiction novel, which is also called *After On*, and which came out last summer. But the podcast acquired a life of its own, and I'm about to publish episode #38 in the series of eight.

As you're about to find out, these first eight episodes have a distinctive format, in that each of them ends with a conversation between me at Tom Merritt - who you might know from CNET, from TechTV, then later from Leo Laporte's network, and now from Tom's own videocast, Daily Tech News Show. In these closing conversations, Tom and I discuss the day's interview, and also a chunk of the book.

The segment you're about to hear will end with one of those conversations. And you may, of course, want to skip it – because while there are no real spoilers, we do talk about the novel, and you may be lost in a few parts. That said, having just relistened to it, I'd say it's a pretty fun conversation. Tom is a very good interviewer, and he interviews me in this part. And we do talk about terrorism, and other topics covered in my interview with Sam. We also talk about tape worms. Plus my own strangely extensive background in the Middle East, my mercifully brief experiences in foster care, and some other random, yet interesting things. So consider listening to that final bit.

And with that - back to my conversation with Sam Harris.

TRANSITION MUSIC

Rob Reid: Do you worry about bio weapons as well?

Sam Harris: Yeah. You just have to imagine weaponizing something akin to the Spanish flu, which have killed something like 50 million people in 1918. Yeah. The sky is the limit there. You could get something that is as easily transmissible and is even more deadly. When you're talking about a bio weapon, the worst possible

case is something that is easily transmissible and it doesn't make you floridly ill for a long enough [crosstalk 01:46:00] do as much damage as you possibly can.

Rob Reid: You sneeze a lot on lots of grapes on lots of people.

Sam Harris: For a good long time before you die.

Rob Reid: Yeah, and then those people are sneezing on grapes and people, and then nobody knows there's an outbreak until there's a million infectees or something like that.

Sam Harris: Yeah. Something like Ebola doesn't have going for it, as bad as it is, as horrible as it is, one of the reasons why it's not scary or is it is very quickly obvious how sick people are. If you're talking about airborne transmission of something that has a very high mortality and a long incubation period, yeah, weaponize that. That is a civilization canceling event if we don't have our-

Rob Reid: And for now, George Church may be the only person who can do it, but in 25 years with biology following what's sometimes called the Carlson curve, which is even steeper than the Moore's law curve, who knows? When 10 people, then 100, then a thousand people. So I'd like to close on something that I wrestle with a lot. You gave a great TED Talk on the risk of super AI. I won't make you replay it here because people can access it. I'll just pull two quotes from to just set the context.

You described the scenario of a super AI having better things to do with our planet and perhaps our atoms than let us continue to have them as being terrifying and likely to occur and also saying it's very difficult to see how they won't destroy us, and I don't think that those are shrill or irrational statements, personally. I also don't think it's really irrational to think that what George Church alone can do today will be the province of many millions of lab techs probably in our lifetimes.

And with those two forces out there, I don't know what scares me more. And I think about proliferating democratizing existentially destructive technology. Just about the only thing I can think of that might protect us against such a thing would be an incredibly benign super AI that has functional omniscience because of its ubiquity in the networks and has functional omnipotence because of its mastery of, who knows, nanotechnology or something else.

But boy, we're both scared about a super AI. It's almost like super AI, can't live with them, can't live without them. How do we navigate those twin perils and do we need to perhaps embrace a super AI as a protective mechanism for democratized super destructive power?

Sam Harris: Yeah. Well, I do think it really isn't a choice. We will develop the most intelligent machines we can build unless something terrible happens to prevent us doing it. So the only reason why we wouldn't build a civilization-

Rob Reid: It gets thrown violently backward.

Sam Harris: Yeah. So, you know, George Church loses his mind or one of his techs does, and we have some pathogen that renders us incapable of keeping our progress going on the technology front. And you just have to imagine how bad that would have to be in order to actually stop the march of progress. [crosstalk 01:49:09] You have to have a world where no one understood how to build a computer again and no one ever understood how to build a computer again, going forward.

Rob Reid: Beyond [cannical 01:49:21] for Liebowitz type of destructiveness. Yeah.

Sam Harris: So if it's not that bad, we will keep making progress, and you don't need Moore's law, you just need some incremental progress to continue.

Rob Reid: You need a passage of time.

Sam Harris: Yeah.

Rob Reid: At some rates, yeah.

Sam Harris: At some point, we will find ourselves in the presence of machines that are smarter than we are because I don't think there's anything magical about the wetware we have in our heads as far as information processing. So the moment you admit that this can be, that what we call a mind can be implemented on another platform, and there's every reason to admit that scientifically now, and I leave questions of consciousness aside. I don't know that consciousness comes along for the ride, necessarily, if you get intelligent machines.

And ironically, the most horrible vision is one of building super intelligent unconscious machines because in the presence of consciousness, at least you could argue, well, if they wipe us out, at the very least, we will have built something more important than we are. We'll have built gods. We will have built minds that can take more pleasure in the beauty of the universe than we can. Who knows how good the universe could be inhabited-

Rob Reid: In their hands, yeah.

Sam Harris: In their hands, right. But if the lights aren't on, if we've built just mere mechanism that is incredibly powerful, that can be goal directed, but for whom there is nothing that is like to be directed towards those goals, that really strikes me as the worst case scenario because then the lights go out if we-

Rob Reid: We go out. So it sounds like you believe that the super AI is inevitable unless [crosstalk 01:50:54]. So our best shot of surviving is to do all we can to make sure the super AI that one day inevitably arises is benign.

Sam Harris: Yeah, is aligned with our interests. Intelligence is the best thing we have. It's our most valuable resource. It's either the source of or the safeguard for everything we care about. Right?

Rob Reid: And there's overwhelming economic incentives for thousands-

Sam Harris: Get immediately rich.

Rob Reid: Intensely smart people, intensely well-capitalized companies to go screaming down that path.

Sam Harris: Yeah. So all of these incentives are aligned to get into the end zone as quickly as possible, and that is not the alignment we need to get into the end zone as safely as possible. And it will always be easier to build the recklessly unsafe version than it will be to take the further step of figuring out how to make this thing safe.

Rob Reid: Yeah.

Sam Harris: So that's what worries me, but I think it is inevitable in some form. And again, I'm not making predictions that we're going to have this in 10 years or 20 years, but I just think at some point, and the human level bit is a bit of a mirage because I think the moment we have something human level, it is superhuman.

Rob Reid: It blows past that.

Sam Harris: Yeah, that's a mirage, and people are imagining somehow that that's a stopping point. It will barely get there, and then we'll stay there for a long time.

Rob Reid: It could only be the case if we are, ourselves, at the absolute summit of cognition, which just defies common sense.

Sam Harris: We just know that's not true. Just take the calculator on your phone. I mean, that's not human level. That is omniscient with respect to arithmetic. Just having the totality of human knowledge instantaneously accessible through the internet, if we hook these things to the internet, it has a memory that is superhuman and an ability to integrate data that superhuman.

So the moment all of these piecemeal cognitive skills cohere in a system that is also able to parse natural language perfectly, you can talk to it and it understands. It does what you want. All of the answers to the questions are no longer like series answer is where they contain howlers every third trial. But they're the most perceptive, best informed, most articulate answers you're

getting from any mind you ever interact with. Right? Once those gains are made, they won't be unmade. It's like chess. It's like once computers were better at chess than people.

Rob Reid: Yeah.

Sam Harris: Now we're in this sort of no man's land, which again, I think will be fairly brief where-

Rob Reid: The centaur.

Sam Harris: Yeah, the combination of a person and a computer is now the best system. But at a certain point, and I'm amazed that anyone doubts this, but at a certain point, I think it will obviously be the case that adding the ape to the equation just adds noise to the equation, and the computers will be better than cyborgs. And once they are, there's no going back from that point.

And it may not be everything. There may be things we neglect to build into our AIs that are turn out to be important for human common sense or this is the scary thing. We don't know what is required to fully align an intelligent system with our wellbeing. And so we could neglect to put something like our common sense, because we don't perfectly understand it, into these systems.

And then you can get errors that are deeply counterintuitive and is analogous to Nick Bostrom's cartoon, thought experiment of the paperclip maximizer. Who would build such a machine? Well, we wouldn't, but we could build a machine that, in the service of some goal that is obviously a good one, could form some instrumental goal that we would never think an intelligent system could form, and that we would never think to explicitly prevent. And yet, this thing is totally antithetical to everything good.

Rob Reid: Yeah, it reaches an equilibrium where it says more paperclips, good. Going to do that for awhile.

Sam Harris: Yeah.

Rob Reid: And soon, the universe has paperclips. Well, Sam, you have been extravagantly generous with your time. I appreciate it.

Sam Harris: Not at all. It's a pleasure.

Rob Reid: And thank you very kindly. We will, I'm sure, remain in touch.

Sam Harris: Yeah. Yeah, and I wish you the best of luck, needless to say, with your book and the podcast and everything else.

Rob Reid: Thank you, kindly.

Sam Harris: It's a great idea that you're combining both in this way. I think obviously this will be the frontier of creative use of these new media. It's great to see you doing it.

TRANSITION MUSIC

Tom Merritt: Rob and I will now discuss the interview and sometimes make specific reference to the novel, *After On*, particularly pages 289 through 380, which are officially on today's roster. So if you're not reading the book, you could be confused by some things we say that you haven't read about, but since you've gotten this far, you'll still probably enjoy the discussion, whereas if you're planning to read the novel, beware of spoilers. Now Rob, the political turmoil in the Middle East has some personal resonance for you, doesn't it?

Rob Reid: Yeah. I actually first went to the Middle East for a period of months when I was 17. I was an exchange student. I grew up in southwestern Connecticut and I became part of this exchange program whose philosophy was, we don't want to just be a travel agency for overprivileged high school kids, so you just tell us that you want to be an exchange student, and then we are going to tell you where you're going.

And I was of the opinion that would be great to go to New Zealand. You could express a preference. I told them that and they said, "Great, you're going to Cairo." So I spent a summer in Cairo when I was 17, and then I got to college and I had been so intrigued by all that I'd learned and seen and done there that I basically majored in Middle Eastern history. Modern Middle Eastern history was technically an international relations major, and I also studied Arabic at great length.

And then when I graduated, I got a Fulbright grant and I went back to Cairo to live and to get really deep into Arabic and to research the political opposition in Egypt. And I wanted to research both the religious and irreligious political opposition, the religious and the secular. But the religious folks wouldn't talk to me. They found it very suspicious that here was this American guy who spoke almost fluent Arabic. It's a lot rustier now, but in those days it was almost fluent. He seems to go to the embassy from time to time, so I gave them the creeps.

But the secular folks did talk to me, and the secular folks were in a hard position because the Mubarak regime was very tough with all opposition, as they're very well known for. And the religious folks were much more ascendant, they were speaking to a much larger percentage of the Egyptian population than the secular opposition. And so the secular people were getting it from Mubarak and the more violent end of the religious opposition was very opposed to them as well.

And a guy named [Furoc Fota 02:01:13] who was pretty prominent in the secular opposition, he was Christian, about 10 percent of Egyptians are, but he also happened to pursue a secular approach to government. He was somebody I

spent a fair amount of time with, and not long after I got home from Cairo, he was assassinated. And that was the beginning of the violent resistance in Egypt, which really hadn't ... There had been very little violence in Egypt since Sadat's assassination, which at that point, was almost a decade in the past.

And that was the beginning of an uptick in violence in Egypt that continues to this day. I've gone back to the Middle East a great number of times since then. I've done work as an election observer, I've done not for profit work, I've been an advisor, a pro bono advisor to lots of startups. So I have a deep history in the region and the issue of terrorism in general. Instability is one that I've been focused on, just in my spare time and with spare cycles for decades.

Tom Merritt: We're going to diverge briefly from the topics you talked about with Sam, but you have another personal connection to the storyline that most readers probably aren't aware of.

Rob Reid: Yeah. So one of the interesting things about Flutter, it's been stated already in the book, is that she is and considers herself to be an orphan. She was, in her case, born of ... Well, I'm not going to go into too much depth of who she was born of, but she considers herself to be an orphan, and that mention has already been made. And as it happens, I was myself born into foster care in New York City.

When I was talking to somebody after I'd finished the first draft of this book and was talking about the principal themes, I mentioned this thing about belonging, and connection to family, and being born to one family and trying to find another family, and so forth, which again, no spoilers. But Flutter is an orphan, so these issues do come up later. And one friend of mine said, "Oh wow, it's so funny that you built in something that is so deeply personal."

And it's like, what? I literally did not see the connection. I saw the connection with the Middle East stuff, I saw the connection with the tech stuff. Those are big parts of my background. That completely missed me. So as we'll get deeper into the book and the second half, these issues will come up. And it is interesting to note that the author was, momentarily for a period of just a very small number of months, technically an orphan, but I got adopted by a wonderful family in Connecticut that is my family and raised me in very stable circumstances.

Tom Merritt: Well, you lucked out.

Rob Reid: I did.

Tom Merritt: For sure.

Rob Reid: I definitely did.

Tom Merritt: But it is interesting how the templates of our consciousness seep into our writing.

Rob Reid: Or flood in and you don't even seem to notice it until somebody points it out two years into the project.

Tom Merritt: Okay. Let's return to the grim topic of militarized nihilism that you and Sam discussed. Tell us a bit more about this [inaudible 02:03:59] movement that figures very prominently in the novel.

Rob Reid: Yeah. So [inaudible 02:04:02] which has already been in the book, is clearly a super nihilistic organization, and it goes beyond any organization that is out there right now in its nihilism. And also, I'll say that all of its precepts run very, very counter to any orthodox or even fringe interpretations of Islam. Unfortunately, that can be said of a lot of violent organizations, that their interpretations of Islam have very little, if any, basis in Islamic scripture, but that is absolutely true of [inaudible 02:04:37].

So they are a work of fiction. That's not me playing my science fiction writer card. That's me playing my horror writer card, but they are based really onto actual historic things. So first of all, the organization has a lot in common with Boko Haram, which rose up in Nigeria and is a really terribly violent, unbelievably brutal organization famous for kidnapping hundreds of school girls and a lot of other just awful stuff.

Tom Merritt: The name literally means western education is prohibited.

Rob Reid: Is prohibited. And [inaudible 02:05:09] as I say in Egyptian Arabic, basically means army of reckoning. And so it's all about the day of reckoning, which is judgment day. It's my fictitious creation. The other thing that [inaudible 02:05:22] is sort of based on, is a person named Anwar al-Awlaki, who was an American citizen but also a very, very radical imam and also Yemeni. He was a famous case some years ago because he was the first American citizen that we deliberately, we the United States, not we Rob and Tom, for those of you who are listening, the United States deliberately targeted for assassination with a drone strike from a predator.

And he is well known because he speaks, I'm sorry, spoke absolutely flawless idiomatic English. And so the sermons that he delivered on terrorism and how it's perfectly reasonable and just to slaughter civilians at home have inspired a number of English speaking radicals such as the person who perpetrated the Fort Hood atrocity, such as the so-called Underwear Bomber, and others because he speaks in this very accessible English.

So [inaudible 02:06:20] is emerged from a terrible military situation, civil war situation in central Africa like Boko Haram, and it also features this charismatic person who infects people with ideas because of his charisma and his

immaculate mastery of language. So that's where [inaudible 02:06:40] comes from.

Tom Merritt: Rob.

Rob Reid: Yes?

Tom Merritt: Thankfully, authors of science fiction and thrillers are under no obligation to write about things they think will actually happen. Right?

Rob Reid: Mm-hmm (affirmative).

Tom Merritt: How worried are you really about synthetic biology terrorism?

Rob Reid: Well, this is going to harken back to some of the stuff that we talked about with our Andy Hessel episode, but that was a few weeks ago. I actually think ... Well first of all, just to recap that, what's called the Carlson curve, which tracks the speed at which synthetic biology is getting better and better and cheaper and cheaper, as I mentioned previously a few weeks ago, that is moving so much faster than Moore's law curve. And we all know how transformative that's been.

It is only a matter of time before lots and lots of people, not just carefully vetted geniuses and high end labs, have the ability to potentially do terrible things with synthetic biology. And unless we freak out about it a lot now and in the coming decades, we could be in a situation in which lots of people will be in a position to hit print and do terrible things. And as I mentioned at the end of the interview with Sam, the bad guys don't have to do 95 percent of the heavy lifting there. Their raw material is that tiny, tiny fraction of incredibly unhappy people who are going to commit suicide this year. And that's about a million people worldwide.

A tiny fraction of them are going to be in a state of mind where they are willing to take as many people with them as possible. And as I mentioned in the interview with Sam, the force multiplier when somebody gets to that place is weaponry. They have mass stabbings in China that kills fewer people than mass shootings in the United States where guns are widely available. If you have an airplane like that Germanwings pilot, Andreas Lubitz or Osama Bin Laden's people, you kill a lot more people than those who do not have airplanes.

When synthetic biology gets to this hyper distributed state, I don't know how we keep an eye on, perhaps, a million people who may do something awful to themselves and a fraction might do something awful to the world if they're in that position. The only way I can think of, and this is scary, the only source of surveillance that could possibly keep a lid on that in my mind, and hopefully this is just a failure of imagination on my part and somebody else, perhaps a listener, will come up with a better idea.

But to me, that would be a good job for a benign super AI, a super AI of the sort that we have talked about in the book, a secret agent Brock Hogan has warned us about, could be as stated in a couple places in the book, functionally omniscient because it would be ubiquitous in the network and it can be functionally omnipotent because we'd probably master things like nanotechnology that lie deep in our future. If that super AI were hearing everything and loved us as much as we would like it to love us, it would be in a position, like Flutter in the book, to derail almost any kind of plot.

So it's almost like super AIs, can't live with them, can't live without them. That is a really, really interesting and important dichotomy to me, and they're both scary paths. I don't know what the third one is where we don't deal with either one of those things. I'd sure like to find it, but that's something that I'm wrestling with as a result of having written this novel, and as a result of having done all the research that went into it, and the following research that you and I are doing for these podcasts.

Tom Merritt: You couldn't just be satisfied with one doomsday scenario.

Rob Reid: Nope, gotta have two.

Tom Merritt: Gotta have two. Which are you more worried about, super AI or syn bio terrorists? Or is it equal?

Rob Reid: I'm worried about the fact that I really have no idea which one I'm more worried about. I think it's something that we have to have some of the brightest people in our society who understand these technologies very, very deeply really thinking very, very hard about. Now the good news is folks like Elon Musk have donated significant amounts of money to look at a super AI safety, and we do actually have some of the brightest people in society when it comes to thinking about where computers might be in five, 10, 15, 20 years starting to think very, very careful, careful about super AI safety.

I'm sure there are others who are thinking about syn bio, but however many there are, it's impossible to have too many people thinking about it. So I guess I'm equally scared about both and what scares me is I usually have a pretty good opinion about which of two things scare me more, and in this case, I don't.

Tom Merritt: Okay. So the destruction of humanity as a squad goal aside, let's try to end on a slightly cheerier note. How about that decrypted file that explains why it's so awesome to use a social network to get everybody to agree to surveillance? Wherever did you get that from?

Rob Reid: Well, you must be talking about Flutter's EULA or the "War and Peace of EULAs" as the Whistleblowing's Blog described it. I actually got that from interviews with Cindy Cohn and some of the people who work at EFF before I started writing the book. So unfortunately, they're not on tape because I was just

scribbling notes back then. I didn't know what was going to write about. But one of the people who worked for her made me aware of a study that I cited actually in an earlier podcast in which it states that it would take all humans, on average, one and a half months per year just to read the privacy agreements of the websites that they use.

Now, I'll add that in most websites, the privacy agreement is shorter than the EULA. And the things that you could get somebody to agree to in a EULA basically has no bounds. Also, many EULAs, as Facebook and others have shown, right in the EULA that it is incumbent upon you to keep track of changes of the EULA because you're bounded not only by that EULA, but by future versions of the EULA. And Facebook does push out these notifications that we all delete without reading, I'm sure, that says, "Hey, we've updated our EULA, our privacy agreement." So it really is probably a full-time job to read that stuff.

So I don't know if that's a lot more cheerful than the super AI or actually, you know, it's more cheerful because that's just like legalese. That's not destruction of human.

Tom Merritt: Relatively speaking.

Rob Reid: So that is an interesting topic, and I'm glad that you raised it.

Tom Merritt: Now, there's a lot of interesting topics. We're so far into the book where there's social shaming, there's decryption by quantum computing. We've talked about that before. But Rob, let's talk about tapeworms, tapeworms being good for you because they keep the immune system from being bored.

Rob Reid: Yes.

Tom Merritt: This is not an essential part of the plot, necessarily.

Rob Reid: Well, it might be. We don't want to spoil anything. People, keep an eye on those tapeworms. It may just be a red herring or red tapeworm. Yes. There are people who were actually pretty smart who believe just that. And so this is sort of the oversanitisation theory. I don't know if it's called that, but let's just call it that. There are folks who have argued that we have so sanitized our world that people grow up without enough exposure to germs and basically become incredibly vulnerable.

We didn't used to have peanut allergies and now peanut allergies are incredibly widespread amongst younger folks, and we don't really know why. But one kind of "fun" explanation might be that we have scrubbed so much stuff that kids just can't handle stuff. Now, a derivative of that theory and I don't know it well enough to utter it articulately, but basically a derivative of that theory has it that if you have a parasite in your system, it's particularly good for toning down autoimmune conditions because autoimmune conditions are basically your own

immune system attacking your system, and it can lead to mild but embarrassing circumstances like dandruff.

That can be an autoimmune condition. It can lead to terrible things like basically this is a retrovirus, but AIDS turns into being an autoimmune condition, and there's lots of them. And so there are folks who think that if your immune system is so underemployed and so bored that it's starting to give you dandruff or create other things that are even worse, maybe you give it a hobby, and a tapeworm could be just such a hobby.

And there was a really good This American Life episode years ago in which the guy was talking actually about hookworm and convinced that it had cured him of some awful condition. So that school of thought is out there, but in the four corners of this novel, obviously that was kind of a comedic scene with these lunatics chasing the latest fad. And I'll just say, there are quite a few health related fads in Silicon Valley that have to deal with crazy diets, that have to do with crazy exercise regimes, that have to do with crazy sleep regimes, that have to deal with all kinds of stuff. And it was a little bit fun to spoof that by talking about our character, Raj, being really fired up about getting a tapeworm.

Tom Merritt: Well, I would like to see more evidence based research confirming this before I'm going to put a tapeworm in my own self.

Rob Reid: I would water rather have that as well, certainly.

Tom Merritt: And I hope everyone considers that as well, but that was really funny. So next week, we finally get to the guts of the novel and we talk about the risk of super AI. We'll be covering pages 380 through 464, and our guest could be a super AI.

Rob Reid: Could be. We're going to keep that a secret until next week. You're going to have to tune in to find out who that guest is.

So Ars Technica listeners - here we conclude the fourth and final installment of my interview with Sam Harris. I do hope you enjoyed it. In case you're interested, the current episode of my podcast is an interview with Great Britain's Astronomer Royal, Martin Rees. Martin and I talk about the most eerie and violent phenomena in the known universe. Specifically, gamma ray bursts in the violent department; and fast radio bursts in the eerie department. We also spend a great deal of time discussing the existential risks society might face in the 21st century..

You can find the Martin Rees episode by visiting my site, at after-on.com. Or, just type the words after-on into your favorite podcast player, and scroll through the episodes. There, you'll find lots of stuff about life sciences - above all, genomics and synthetic biology. Conversations about robotics, privacy and government hacking, cryptocurrency, astrophysics, drones, and a whole lot more.

Or, you could just join me next week, here on Ars.