

SAM HARRIS

THE BLOG

Complexity & Stupidity

Podcast Transcript

[Ethics](#) | [Neuroscience](#) | [Philosophy](#) | [Physics](#) | [Science](#) | November 13, 2016



(Photo via the [Anguskirk](#))

In this episode of the Waking Up podcast, Sam Harris talks to biologist David Krakauer about information, complex systems, and the future of humanity.

David Krakauer is President and William H. Miller Professor of Complex Systems at the Santa Fe Institute. His research

explores the evolution of intelligence on earth. This includes studying the evolution of genetic, neural, linguistic, social and cultural mechanisms supporting memory and information processing, and exploring their generalities. He served as the founding Director of the Wisconsin Institute for Discovery, the Co-Director of the Center for Complexity and Collective Computation, and was Professor of mathematical genetics at the University of Wisconsin, Madison. He previously served as chair of the faculty and a resident professor and external professor at the Santa Fe Institute. He has also been a visiting fellow at the Genomics Frontiers Institute at the University of Pennsylvania, a Sage Fellow at the Sage Center for the Study of the Mind at the University of Santa Barbara, a long-term Fellow of the Institute for Advanced Study in Princeton, and visiting Professor of Evolution at Princeton University. In 2012 Dr. Krakauer was included in the *Wired* Magazine Smart List as one of 50 people “who will change the World.”

For information about the Santa Fe Institute: www.santafe.edu

The article discussed in this podcast: [The Empty Brain](#)

* * *

Welcome to the Waking Up podcast. This is Sam Harris. Today, I'll be speaking with David Krakauer, who runs the Santa Fe Institute, one of the most interesting organizations, scientifically, you'll find anywhere. David is a mathematical biologist. He has a PhD in evolutionary theory from Oxford, but being at the Santa Fe Institute puts him at the crossroads of many different areas of inquiry. We'll talk a little bit about what the institute is, but given that its focus is on complex systems, the people there attempt to understand complexity using every scientific and intellectual tool available.

David knows a lot about many things, as you'll hear in this conversation. We'll start by covering some foundational concepts in science—like information, complexity and intelligence—and move on to their implications for society and culture in the future. I loved talking to David, and I hope you enjoy the ground we cover.

Harris: David, thanks for joining me on the podcast.

Krakauer: Pleasure to be with you.

Harris: You gave a fascinating lecture in Los Angeles that I want to talk about. I'd like you to track through that as much as you can without your visuals. I'm interested in the importance of culture, especially the artifacts that we create to support human intelligence, and in resisting our slide into stupidity, which was the focus of your talk. But before we get there, let's set the stage a bit. Tell us about your scientific interests and background.

Krakauer: My scientific interests, as I've come to understand them, are essentially grappling with the evolution of intelligence and stupidity on earth. It's quite common for people to talk about intelligence. It's less common for people to talk about stupidity, even though, arguably, it's more common. My background is in mathematical evolutionary theory, and I work on information and computation in nature. That includes the nature that we've created, that we call technology, and where it came from, what it's doing today, and where it's going in the future.

Harris: Now you're running the Santa Fe Institute. Its existence seems to be predicated on the porousness of these

boundaries between disciplines, or even their nonexistence. Describe the institute for people who are not familiar with it.

Krakauer: The Santa Fe Institute is in Santa Fe, New Mexico, as the name would suggest. It was founded in the mid '80s by a group of Nobel laureates from physics and economics and others who were interested in trying to do for the complex world what mathematical physics had done so successfully for the simple world. I should explain that. The simple world would be the solar system, or inorganic chemistry, or black holes. They're not easy to understand, but you can encapsulate their fundamental properties by writing down a system of equations. The complex world, which basically means networked adaptive systems, could be a brain, a network of neurons. It could be a society; it could even be the Internet. In those networked adaptive systems, complex systems, the kinds of formalisms that we had created historically to deal with simple systems failed.

That's why we don't have Maxwell's equations for the brain, right? We have large textbooks with many anatomical descriptions, some schematic representations of function, and some very specialized models, and the question for us to justify is, Are there general principles that span the economy, brains, the Internet, and so on, and what is the most natural way of articulating them mathematically and computationally?

Harris: How is SFI different from the Institute for Advanced Study at Princeton, where I think you also were at one point?

Krakauer: Yes, that's right. The IAS in Princeton is a lot older. It was founded in the '30s, we were founded in the '80s. IAS is an extraordinary place, but the model, if you like, is much more traditional. IAS has tenure, it has departments, and it has schools. We do not have tenure, we do not have departments, and we do not have schools. In some sense, they've replicated a very successful model that is the university model. We decided to start from a blank slate, and we asked the question, "If you were reinventing a research institute based on everything that we now know, post-scientific revolution, post-technological revolution, etc., what should it look like?" So it's a more radical model, and we decided very early just to discard any mention of disciplines and departments and focus as hard as we could on the common denominators of the complex systems that we were studying.

Harris: And it's truly interdisciplinary. You have economists and mathematicians and biologists and physicists all throwing in their two cents on the same problems. Is that correct?

Krakauer: Absolutely. You know, there's all this debate now about the demise of the humanities. But we, from the very beginning, decided that that wasn't a worthwhile distinction—between the natural sciences and the humanities. So we've been working on the archaeology of the Southwest and using computational and physical models since the '80s, and we've produced what is by now a very well known series of theories for why, for example, some of the native civilizations of the American Southwest, the origin of ancient cities, declined. All of these are based on computational and energetic theories and close collaborations between archaeologists and, say, physicists. I don't like to call the way we do it interdisciplinary, because that's in some sense genuflecting in the direction of a superstition that I know people take seriously. So what happens when you ignore all of that and say, "Let's certainly use the skills that we've acquired in the disciplines, but let's leave them at the door and just be intelligent about complex problems"?

Harris: What you have is an institutional argument, it seems to me, for the unity of knowledge, or *consilience*. The boundaries between disciplines are much more a matter of university architecture and the bandwidth limits of any individual

career, wherein it takes a long time to get very good at one thing. So by definition, someone starts out in one area as opposed to another and spends rather a long time there in order to get competent. So I think what you're doing there is very exciting.

Krakauer: Thank you.

Harris: Before we get into your talk, I want you to enlighten me and our audience about a few things, because you are going to use some concepts that I think are difficult to get one's head around. The first is the concept of information. There are many senses in which we use this term, and not all of them are commensurable. It seems to me that there is a root concept, however, that potentially unites fields like genetics and brain science and computer science and even physics. How do you think about information?

Krakauer: We've talked about this before, Sam. It's sometimes what I call the *m-cubed mayhem*. That is m raised to the power of three mayhem. The mayhem comes from not understanding the difference between mathematics, the first m , mathematical models, the second m , and metaphors, the third. And there are terms—scientific terms, mathematical terms—that are also used idiomatically or have a colloquial meaning, and they very often get us into deep water: energy, fitness, utility, capacity, information, computation. We all use them in our daily lives, probably very effectively, but they also have a technical meaning. And what happens often is that arguments flare up because one person is using a term mathematically and another person metaphorically, and they don't realize they're doing this. I don't mean to say that there is only a mathematical definition of information, but it's worth bearing in mind that when I talk about it, that's what I mean. So that's the first point. It has a beautiful, scientific, storied history, starting with essentially the birth of the field that we now call statistical mechanics. This was essentially Boltzmann trying to understand the arrow of time in the physical world, the origin of irreversibility. You know, why is it that you can crack and break an egg, but the reverse almost never happens? Why is it that you can burn wood into ash and smoke, but the reverse almost never happens?

He created, in the 1870s, a theory called the H-Theorem, where he essentially had in mind lots of little billiard balls bumping into each other chaotically. You start with a fairly ordered billiard table, but at the end, through repeated collisions, they're distributed rather randomly all over the table. That was Boltzmann. He thought maybe the underlying molecular structure of matter was like lots of little billiard balls, and the reason why we observe certain phenomena in nature as irreversible is because of molecular chaos. That was formalized later by a very famous American physicist, Josiah Willard Gibbs. But many years later, the idea was picked up by an engineer working at Bell Labs, Claude Shannon. He realized that there was a connection between physics and irreversibility—an arrow of time and information. It was a very deep insight.

And before he explained how that worked, what did Claude Shannon do? He said, "Look, here is what information is. Let's say I wanted to navigate from one part of the city to another, from A to B. In a car, I could just drive around randomly. It would take an awful long time to get there, but I might eventually get there. Alternatively, I could give you a map or driving directions, and you'd get there very efficiently. The difference between the time taken to get there randomly and the time taken to get there with directions is a measure of information." Shannon mathematized that concept and said, "That is the reduction of uncertainty. You start out not knowing where to go, you get information in the form of a map or driving directions, and then you get there directly." He formalized that, and he called it information.

It's the opposite of what Boltzmann and Gibbs were talking about. It's a system. Instead of going from the ordered into the disordered state, the billiard balls on the table start maybe in a lattice and end up randomly distributed. It's going from a

state of being random, because you don't know where to go, to becoming ordered. It turns out that Shannon realized that information is in fact the negative of thermodynamic entropy, and it was a beautiful connection that he made between what we now think of as the science of information and what was the science of statistical physics.

Harris: Let's bring this into the domain of biology, because I've been hearing with increasing frequency the idea that biological systems and even brains do not process information, and that the analogy of the brain to a computer is no more valid than the analogy of it to a system of hydraulic pumps, or wheelworks powered by springs and gears, or a telegraph. As you know, these were all old analogies to the most current technology of the time. But there was an article in *Aion* magazine that probably a dozen people sent to me, which made this case very badly. And you and I spoke about this briefly when we first met. No one, to my knowledge, thinks that the brain is a computer in exactly the way our current computers are computers. We are not talking about Von Neumann architecture in our brains.

But the idea that the brain doesn't process information at all, and that to claim that it does is just as crazy as claiming that it's a mechanism of gears and springs, strikes me as fairly delusional. However, I keep meeting people who will argue this, and some of them have careers in science. So I was hoping we could talk a little bit about the ways in which biological systems, in particular brains, encode and transmit information.

Krakauer: This takes me right back to my m-cubed mayhem, because that's a beautiful example in that paper of the author not knowing the difference between a mathematical model and a metaphor. You talk about springs and levers and their physical artifacts, right? And then there are mathematical models of springs and levers, which are actually used in understanding string theory. So let's talk a bit about the computer and the brain. You mentioned Von Neumann. It spans elegantly that spectrum from mathematics to mathematical models to metaphors. The first real theory of computing that we have is due to Allan Turing in the 1930s, and he was a mathematician.

Many people know him from the movie *The Imitation Game* and for his extraordinary work on Enigma and decoding German submarine codes in the Second World War. But what he's most famous for in our world is answering a really deep mathematical question that was posed by the German mathematician David Hilbert in 1928. Hilbert said, "Could I give a machine a mathematical question or proposition, and it would tell me in a reasonable amount of time whether it was true or false?" That's the question he posed. Could we in some sense automate mathematics? And in 1936, Turing, in answering that question, invented a mathematical model that we now know as the Turing machine, and it's a beautiful thing. I'm sure you've talked about it on your show before. Turing did something remarkable. He said, "You know, you can't answer that question. There are certain mathematical statements that are fundamentally uncomputable. You could never answer them." It was a really profound breakthrough in mathematics when he said there are certain things in the world that we could never know through computation. Years later, Turing himself, in the '40s, realized that in solving a mathematical problem, he had actually invented a mathematical model, the Turing machine. And he realized the Turing machine was actually not just a model for solving math problems; it was actually the model of problem-solving itself. And the model of problem-solving itself is what we mean by computation. Then, in the 1950s, actually '58, John Von Neumann wrote a book, the famous book *The Computer and the Brain*.

They said perhaps what Alan Turing did in his paper on intelligent machinery has given us the mathematical machinery for understanding the brain itself. At that point, it became a metaphor. John Von Neumann himself realized it was a metaphor, but he thought it was very powerful. So that's the history. Now, back into the present. As you point out, there is a tendency

to be a bit, you know, epistemologically narcissistic. We tend to use whatever current model we use and project that onto the natural world as almost the best-fitting template for how it operates.

Here is the value, or the utility and disutility, of the concept. The value of what Turing and Von Neumann did was to give us a framework for starting to understand how a problem-solving machine could operate. We didn't really have in our mind's eye an understanding of how that could work, and they gave us a model of how it could work. For many reasons, some of which you've mentioned, the model is highly imperfect. Computers are not robust. If I stick a pencil in your CPU, your machine will stop working. But I can sever the two hemispheres of the brain, and you can still function. You're very efficient. Your brain consumes about 20% of the energy of your body, which is about 20 watts. It's 20% of a lightbulb. Your laptop consumes about that, and has, you know, some tiny fraction of your power. And they're highly connected. The neurons are densely wired, whereas that's not true of computer circuits, which are only locally wired. Most important, the brain is constantly rewiring and adapting based on inputs, and your computer is not.

So we know the ways in which it's not the same. But as I say, it's useful as a full experiment for how the brain might operate. That's the computer term. Now let's take the information term. That magazine article you mentioned is criticizing the information concept, not the computer concept—which is limited, and we all agree, but the information concept is not, right? So we've already determined what information is mathematically. It's the reduction of uncertainty. Think about your visual system: When you open your eyes in the morning and you don't know what's out there in the world, electromagnetic energy, which is transduced by photoreceptors in your retina and then transmitted through the visual cortex, allows you to know something about the world that you did not know before.

It's like going from the billiard balls all over the table to the billiard balls in a particular configuration. Very formally speaking, you have reduced the uncertainty about the world. You've increased information, and it turns out you can measure that mathematically. The extent to which that's useful is proved by neuro-prosthetics. The information theory of the brain allows us to build cochlear implants. It allows us to control robotic limbs with our brains. So it's not a metaphor. It's a deep mathematical principle. It's a principle that allows us to understand how brains operate and reengineer it. I think the article is so utterly confused that it's almost not worth attending to.

Now, that's information. Information processing: If that's synonymous in your vocabulary with computing in the Turing sense, then you and I just agreed that it's not right. But if information processing is what you do with Shannon information, for example, to transduce electromagnetic impulses into electrical firing patterns in the brain, then it's absolutely applicable—and how you store it, and how you combine information sources. When I see an orange, it's orange color, and it's also a sphere. I have tactile, mechanical impulses, and I have visual electromagnetic impulses. In my brain, they're combined into a coherent representation of an object in the world. The coherent representation is in the form of an informational language of spiking. It's extraordinarily useful. It has allowed us to engineer biologically mimetic architectures, and it's made a huge difference in the lives of many individuals who have been born with severe disabilities. So I think we can take that article and shred it.

Harris: As I was reading the article, I was also thinking of things like genes that can be on or off. There is a digital component going all the way down into the genome, and the genome itself is a kind of memory, right? It's a memory for structure and physiology and even for certain behaviors that have proved adaptive in the past. It's a template for producing

those features in future organisms.

Krakauer: That's exactly right. That's the great power of mathematical concepts. Again, we have to be clear in making distinctions between the metaphor of memory and the mathematical model of memory. The beautiful thing—that's why mathematics is so extraordinarily powerful—is that once we move to the mathematical model of memory, exactly as you say, you can demonstrate that there are memories stored in genes, there are memories stored in the brain, and they bear an extraordinary family resemblance through the resemblance in the mathematical equations. You described it as "consilience," in Ed Wilson's term. You could describe it as unification in the language of physics.

Where we run into trouble is if we don't move to mathematics but remain in the world of metaphor. There, of course, everyone has a slightly different matrix of associations, and you can never fully resolve the ambiguities.

Harris: Let's forget about the math for second and talk about something that's perilously close to metaphor. We are talking about cause-and-effect relationships that, in this case, reliably link inputs and outputs. Even in that article, he was talking about the nervous system being changed by experience—he just didn't want to talk about the resulting changes in terms of "memory" or "information storage" or "encoding" or anything else that suggested an analogy to a computer. But change in a physical structure can produce reliable changes in its capacities going forward. Whether we want to call that memory or learning, or not, physically, that's what we're talking about.

Krakauer: Absolutely, that's what we are talking about. You're right. That's the point. It has to do with this legitimate fear of anthropomorphism, and I think that what we do in these sorts of more exact sciences is try to pin down our definitions so as to eliminate some of the ambiguities. They never go away entirely, but my suspicion, Sam, is that the author of that article will simply find a language that doesn't have its roots in the world of information, and apply these new terms. But we would realize, if we read it through thoroughly, that they were, in fact, just synonyms. He would find himself having to use these terms because they are, to the best of our knowledge, the best terms we have to explain the regularities we observe.

Harris: And yet we don't have to use terms like "hydraulic pumps" or "four humors." We can grant that there have been bad analogies in the past where the details were not conserved going forward.

Krakauer: If you're talking about your cardiac system or your urogenital system, it is entirely appropriate to use Harvey's model, which was the pump, right? The ones that worked have stuck, and I think time will tell whether use of the informational concept will be an anachronism or will have enduring value.

Harris: For those of you who want to read this paper that we've been trashing, I will put the link on my blog beneath where I embed this podcast.

Now, moving on to your core area of interest, David: We've dealt with information. What is complexity?

Krakauer: Yes. That's a wonderful example of these terms that we use in daily life but that also have mathematical meaning. The simplest way to think about complexity is as follows. Imagine you have a very regular object, like a cube. You could express it just by describing its linear dimensions, and that would tell you what a cube is. And imagine you want to explain something at the other end of the spectrum, like a gas in a room. You could articulate that very reliably by just

giving the mean velocities of particles in air. So these two extremes—the very regular, a cube, to the very random, a gas—permit of a description, which is very short. Over the phone or over Skype, I could describe to you very reliably a regular object or a very irregular object.

But now let's imagine you said, "Can you please describe a mouse to me, David?" And I said, "Well, it's a sort of weird tubular thing, and it's got hairs at one end, it's got this long appendage at the other, and etc." It would take an awfully long time to describe. Complexity is essentially proportional to that description. So that's a metaphor. It turns out that mathematically, the complex phenomena live somewhere between the regular and the random. Their hallmark signature is that their mathematical descriptions are long, and that's what has made complex science so hard.

Einstein could write down a beautiful equation like $E=mc^2$ that captures the equivalence between energy and mass and has all these beautiful implications for special relativity in less than a line. But how would you write down an equation for a mouse, which seems like a much more boring thing than energy and matter? You can't. So that's one way, an intuitive way of thinking about a complex phenomenon: How long does the description have to be to reliably capture much of what you consider interesting about it? One point to make immediately is that physical phenomena started off long too. Before Kepler revolutionized our understanding of celestial mechanics, we had armillary spheres with all these epicycles and deferents explaining—incorrectly—the circular motion of a celestial mass. It took a while for us to realize that there was a very compact, elegant way of describing them. And it could be that for many complex phenomena, there is a very elegant, compact way of describing them. For many others, I don't think that will be the case.

So complexity is, as I said, these networked adaptive systems. Complexity itself, as a concept, mathematically tries to capture how hard it is to describe a phenomenon. And as they get more complex, these descriptions get longer and longer.

Harris: You said something about randomness there that caught my ear. My understanding is that randomness generally can't be expressed simply. If I gave you a truly random string of digits, unless you're talking about some method by which to produce it algorithmically—the decimal expansion of pi, for instance, can be compressed—but if it's a truly random series of digits, that's not compressible, right?

Krakauer: That's absolutely right. That's a very important distinction. I can describe the process of generating heads and tails by describing the dynamics of the coin, and that's very short, right? But if I was trying to describe the thing I observe, it would be incompressible, and the description would be as long as the sequence described. In all these cases, we're talking about the underlying causal process that generates the pattern, not the pattern itself. And that's a very important distinction.

Harris: This is the first time I've ever conducted an interview like this, just stepping through definitions, but I think it's warranted in this case. So what is intelligence, and how is it related to complexity?

Krakauer: Intelligence is, as I say to people, one of the topics about which we have been most stupid. All our definitions of intelligence are based on measurements that can only be applied to humans—by and large, humans that speak English or what have you. An IQ test is not interesting if you're trying to calculate the intelligence of an octopus—which I would like to know, because I believe in evolution. I think we need to understand where these things come from, and having a definition that applies just to one particular species doesn't help us. We've talked about entropy and computation, and

they're going to be the keys to understanding intelligence.

Let's go back to randomness. The example I like to give is Rubik's cube, because it's a beautiful little mental model, a metaphor. If I gave you a cube and asked you to solve it, and you just randomly manipulated it, since it has on the order of 10 quintillion solutions, which is a very large number, if you were immortal, you would eventually solve it. But it would take a lifetime of several universes to do so. That is random performance. Stupid performance is if you took just one face of the cube and manipulated that one face and rotated it forever. As everyone knows, if you did that, you would never solve the cube. It would be an infinite process that would never be resolved. That, in my definition, would be stupid. It is significantly worse than chance.

Now let's take someone who has learned how to manipulate a cube and is familiar with various rules that allow you, from any initial configuration, to solve the cube in 20 minutes or less. That is intelligent behavior, significantly better than chance. This sounds a little counterintuitive, perhaps, until you realize that's how we use the word in our daily lives. If I sat down with an extraordinary mathematician and I said, "I can't solve that equation," and he said, "Well, no, it's easy. Here, this is what you do," I'd look at it and I'd say, "Oh, yes, it is easy. You made that look easy." That's what we mean when we say someone is smart. They make things look easy.

If, on the other hand, I sat down with someone who was incapable, and he just kept dividing by two, for whatever reason, I would say, "What on earth are you doing? What a stupid thing to do. You'll never solve the problem that way."

So that is what we mean by intelligence. It's the thing we do that ensures that the problem is efficiently solved and in a way that makes it appear effortless. And stupidity is a set of rules that we use to ensure that the problem will be solved in longer than chance or never and is nevertheless pursued with alacrity and enthusiasm.

Harris: Now we're getting closer to the substance of that lecture you gave. I want you to recapitulate part of it here, because I found it fascinating. In particular, I'm interested in the boundary you drew between biology and culture and the way in which culture is a machine for increasing our intelligence. You also fear that we are producing culture in a way that might be making us biologically or personally less intelligent—perhaps to a dangerous degree. If you could just take us there...

Krakauer: This is a lengthy narrative. I'm going to try to compress it. I'll make it the least complex possible. Most of us are brainwashed to believe that we are born with a certain innate intelligence and we learn things to solve problems, but our intelligence goes basically unchanged. You hear this all the time in conversations. They'll say, "That person is really smart. It's just they never worked very hard and they didn't learn very much. Whereas that person is not very smart, but they learned a great deal, and it makes them look smart." That sort of thing. I think that's rubbish. I think there's a very real sense in which education and learning make you smarter. So that's my premise.

Harris: Let's just pause there for a second. You wouldn't dispute, though, that there are differences in what psychologists have come to call "g" or "general intelligence" and that this is somehow not necessarily predicated upon acquiring new information?

Krakauer:

I would dispute that.

Harris: So you think the concept of IQ is useless, not just for an octopus but in people?

Krakauer: More or less, and I should explain why. I think a lot of recent research is required to understand why. Let's take a canonical example—the young Mozart. People say, “Well, look. Wait a minute. This is a kid who at the age of seven had absolute pitch, and in his teens, you could play him a symphony that he could recollect note for note and reproduce on a score. Surely this is an individual who is born...” What we now understand, of course, is that his father was a tyrant who drilled him and his sister from an extraordinarily young age in acquiring perfect pitch, in the subtleties of musical notation. Consequently, he was able to acquire very young characteristics that normally you wouldn't acquire later, because normally you wouldn't be drilled.

More and more studies indicate that if you subject individuals to deliberative practice regimes, they can acquire skills that seem almost extraordinary. Let's take “g” and the IQ in general. We now know that what it really seems to be measuring is working memory, and many working-memory tasks are correlated; they live in this low-dimensional space that we call “g.” One of the classic studies was the number of numbers that you could hold in your head. I recite a series of numbers, and I ask you to remember them, and 10 minutes later, I ask you. You're not allowed to write them down, but what you do is replay them in your mind. People could do 10, maybe they could do 11, and this was considered to be some upper limit on our short-term memory for numbers.

And yet, in a series of experiments, through very intelligent, ingenious means of encoding numbers, people could remember up to 300. These were individuals, by the way, who at no point in their lives ever showed any particularly extraordinary memory capacity. So the evidence is on the side of plasticity, not innate aptitude. And to the extent that IQ is fundamentally measuring working memory, we now know how to start extending it. That's an important point. I wouldn't deny that there are innate variations. I mean, I am not six-foot-five. I'm not even six-foot, so I will never be a basketball player. So there are functions in the world that are responsive to variation that looks somewhat inflexible. But in the world of the brain, given that it is not a computer, and the wiring diagram is not fixed in the factory but actually adapts to inputs, there is much more hope—and in fact evidence—that the variation is much greater than we had thought.

Harris: So the plasticity and trainability ride atop innate variation. You can have differences in aptitude with and without training.

Krakauer: That's exactly right. I think the open question for us is, How much of that, if you like, innate Lego material is universal? How many of those pieces were preassembled into little castles and cars, which we could then build upon? Whether that, or some people arrive on the stage with an advantage, is actually not known. All I'm reporting is that the current deliberative practice data suggests that's less true than we thought it was.

Harris: Which puts the onus, to an even greater degree than most people would expect, on culture and the rest of the machinery that is outside any individual brain but which is, in a material sense, augmenting its intelligence. So take us in that direction.

Krakauer: Yes. That's a very important point. That's why that connection is important to make. Now, we've basically

understood what intelligence is, what stupidity is, and we understand that we are flexible to an extraordinary degree. Maybe not infinitely so, and as you point out, but the inputs then become much more important than we thought in the past. Now let's move into intelligent—or what sometimes gets called cognitive—artifacts. Here's an example. Your ability to do mathematics or perform mathematical reasoning is not something you were born with. You did not invent numbers, you did not invent geometry or topology or calculus or number theory or anything else, for that matter. They were all given to you if you chose to study mathematics in a class. And what those things allow you to do is solve problems that other people cannot solve.

Numbers are in some sense the lowest-hanging fruit in our mathematical education. So let's look at numbers. There are many number systems in the world. There are very ancient Sumerian cuneiform numbers about 5,000 years old, some ancient Egyptian numbers. And here is a good example of stupidity in culture. Western Europe, for 1,500 years, used Roman numerals, from about the second century B.C. to 1500 A.D., toward the end of the Holy Roman Empire. Roman numbers are good at measuring magnitude, the number of objects, but terrible for performing calculations. What's $X + V$? You know. What's XII multiplied by IV? It just doesn't work, and yet for 1,500 years, the human brain opted to deliberate over arithmetic operations using Roman numerals that don't work. The consequence of that was that for much of their history, Europeans could not divide and multiply. It's extraordinary, because it's unbelievably stupid when you realize that in India and Arabia, they had a number system.

It started in India and then moved to Arabia. It was available from about the second century, and that is the system we use today, which can effortlessly multiply and divide numbers. That's a beautiful example of the interface between culture and our own reasoning. The reason it's so intriguing is because once I've taught you a number system, like the Indian Arabic number system, base 10 number system, you don't need the world anymore. You don't need paper anymore to write it down. You can do these operations in your mind's eye, and that's what makes them so fascinating. I call objects like that, which were invented over the course of centuries by many, many minds, *complementary cognitive artifacts*. Their unique characteristic is, not only do they augment your ability to reason in the form, for example, of multiplying and dividing, but when I take them away from you, you have in your mind a trace of their attributes that you can deploy. It's interesting. That's probably what's new in thinking about the evolution of cultural intelligence.

For a long time, psychologists, cognitive scientists, archaeologists, have understood that there are objects in the world that allow us to do things you couldn't do otherwise. A fork, or a scythe, or a wheel. But there is a special kind of object in the world that not only does what the wheel and the scythe and the fork do, but also changes the wiring of your brain so that you can build in your mind a virtual fork, or a virtual scythe, or a virtual wheel. That, I would claim, is the unique characteristic of human evolution.

Harris: Wouldn't you put language itself into this category?

Krakauer: Absolutely. The reason I separate them is that many people erroneously assumed, up until quite recently, that mathematical reasoning depended on linguistic reasoning, and was in fact just a special form of it. We now know that's not true, and that both humans and nonhuman primates are capable of representing number equally well. In fact, when humans perform mathematics, they are using not the linguistic parts of their brain but the parts that represent number, which we

share with nonhuman primates.

Harris: What else would you put on this list of complementary cognitive artifacts?

Krakauer: The other example that I'm very enamored of is the abacus. The abacus is a device for doing arithmetic in the world with our hands and eyes. But expert abacus users no longer have to use the physical abacus. They actually create a virtual abacus in the visual cortex. And that's particularly interesting, because a novice abacus user like me or you thinks about them either verbally or in terms of our frontal cortex. But as you get better and better, the place in the brain where the abacus is represented shifts, from language-like areas to visual, spatial areas in the brain. It really is a beautiful example of an object in the world restructuring the brain to perform a task efficiently—in other words, by my definition, intelligently.

Maps are another beautiful example of this. Let's imagine we don't know how to get around a city. Over the course of centuries or decades or years, many people contribute to the drawing of a very accurate map. But if you sit down and pore over it, you can memorize the whole damn thing. And you now have in your mind's eye what it took thousands of people thousands of years to construct. You've changed the internal wiring of your brain, in a very real sense, to encode spatial relations in the world that you could never have directly experienced. That's a beautiful complementary cognitive artifact. And then some mechanical instruments: You could say that as you become more and more familiar with an armillary sphere or an astrolabe or a sextant or a quadrant, you have to use it less and less. So you build a kind of a simulation in your brain of the physical object. And at some point, in some cases, you can dispense with the object altogether.

Harris: The other shoe drops: There is another kind of cognitive artifact that you want to talk about. Tell us about the downside to all our cultural creativity.

Krakauer: There is another kind of cognitive artifact. Consider a mechanical calculator or a digital calculator on your computer. It augments your intelligence in the presence of the device. So my phone and I together are really smart, right? But if you take that away, you're certainly no better than you were before, and you are probably worse, because you probably forgot how to do long division, because you're now dependent on your phone to do it for you.

Now, I'm not making a normative recommendation here. I'm not saying we should take people's phones away and force them to do long division. I'm simply pointing out there is a difference. And the difference is that what I call *competitive cognitive artifacts* don't so much amplify human representational ability as replace it. Another example that everyone is very enamored of now, rightly, is machine learning. We have this beautiful example recently of AlphaGo, a deep learning neural network being trained to beat an extraordinary ninth-dan Go player. That machine is basically opaque, even to its designers, and it *replaces* our ability to reason about the game. It doesn't augment it.

Another example would be the automobile. This is one of my favorites, because automobiles clearly allow us to move very quickly over an even surface. And we are utterly dependent on them, especially here in the Southwest, where I live. But if you took my car away, I would be no better than I was before, and probably I would be worse, because I would be unfit. I had been so accustomed to sitting in the car for a long time. Moreover, it's a dangerous artifact, because it kills so many people. So the car is a beautiful example of a competitive cognitive artifact that we have accepted, because its utility value is

so high, even though it actually compromises our ability to function without it.

I think the world can be divided into these two kinds of cultural objects. And the question, of course, is, Can we depend on these objects always being around? In the case of competitive cognitive artifacts, if we cannot, then we should worry, right? Because when they're taken away, we'll probably be worse off than we were before.

Harris: The car is an interesting example, because it's just about to make the next iterative leap into being even more competitive, with the self-driving car. You can easily envision a time when self-driving cars are the norm, because they'll be so much safer than ape-driven cars, and yet, that will almost certainly be a time when people's driving skills will have deteriorated virtually to the point of nonexistence. We won't be able to take over the controls in any competent way, once we've lived in the presence of this technology long enough.

Krakauer: That's actually right, and it's interesting, because the driverless car does several things at once. It eliminates the leg, and it eliminates our mapmaking ability. So it actually assaults several cognitive capacities at once. And I really think that the debate that we need to be having—and this is where I've been somewhat frustrated by all of the singularity debate, or the AI doom and gloom debate, because the argument that seems to be playing out in tech circles is, Will we create a machine that will turn around and say, "You expend too much energy. You have a disrespect for the environment. I'm gonna make you a battery." The Matrix nightmare. Whereas the real discussion we should be having, the imminent and practical debate, is what to do about competitive cognitive artifacts that are already leaving an impression on our brains that is arguably negative.

When I have a discussion with somebody about this topic, as a rule, the only recourse they have—and it's totally reasonable—is they're not going away. But there is something else that hasn't been mentioned here, and this is really interesting, which has to do with the complex system of the brain and the domino-like effect and interconnectedness of representation systems. For example, it's been known for a long time that if you become competent at the abacus, you're not just competent at arithmetic. It actually has really interesting indirect effects on linguistic competence and geometric reasoning. It doesn't have a firewall around it such that its functional advantages are confined to arithmetic. And in fact, I think that's generally true for all interesting complementary cognitive artifacts.

So if I give you a fork or chopsticks or a knife, it's true that you're better able to manipulate and eat your food, but you also develop dexterity, and that dexterity can be generalized to new instances. And for me, the main concern is not only that the world will go south and we'll no longer have highways and cars, but also the indirect, diffusive impact of eliminating a complementary cognitive artifact, like a map, on other characteristics we have. Your familiarity with mapmaking and topographical, topological, geometric reasoning is generally valuable in your life, not just in navigating across the city. So taking away a map doesn't just make you worse at getting from one door to another, it makes you worse in many ways. I would strongly claim that this is where the debate needs to be had, because I don't have an answer.

Harris: I think there are probably many other examples of this. I'm not very close to this research, but I know that many learning experts believe that cursive writing, for instance, is important to learn—even though we're living in an increasingly type-written world, which will soon be a voice-recognition world—because it's intimately connected with the acquisition of literacy itself. The pace at which one writes cursively is apparently important. And the physical linking of letters is not

surrounded by a firewall. It's actually related to learning to read well.

Krakauer: A good example of this, which both Einstein and Frank Lloyd Wright depended upon, was wooden cubes. Early in their youth, they both became very enamored of these cubes and would construct worlds out of cubes, like Minecraft. And both of them claimed, Frank Lloyd Wright in the case of architecture and Einstein in the case of the geometry of the universe, that the intuitions they built up playing with these cubes were instrumental in their later lives. I would claim the same is true for maps. If you know how to navigate through a true space, like a Euclidean space or a curved space on the surface of the earth, that allows you to think about different kinds of spaces, relationship spaces, idea spaces. The notion of a path from one idea to another, as a metaphor, actually has an immediate and natural implementation in terms of a path in real space. You can see immediately how these things are of value more broadly.

Harris: You just said a moment ago that you weren't making any normative claims, but the norms just come flooding in once you begin talking about the possible changes in our cognition, and perhaps even in our ethics, once we begin to change the cultural landscape with competitive as opposed to complementary technology. So let's talk about the kind of normative claims one might want to make here.

Most of us want to maximize our capacity to get what we want out of life. And if we were convinced that some technology was reliably diminishing our individual abilities, or producing a spectrum of negative effects that we had not considered, once this came to our attention, we might want to make a change. Then there are also collective norms, where we talk about whole societies being capable of a certain kind of creativity or cooperation, whereas other societies are not. Some societies are in a perpetual state of self-siege or civil war. So how do you think about individual and collective norms in this context?

Krakauer: It's very tricky. The first thing I should say is, I do agree with you that there are, in some domains, absolutely better ways of being. I'll give you an example from writing code for computers. Imagine that we still had to write with punch cards—there would be no word processor. The idea of taking a typewriter and connecting it to a computer was an extraordinary invention—and later, word processors and everything else. Well, let's go a little further. Let's imagine that you could only interact with the computer using machine code or binary. There would be no software as we understand it today, because the projects would always be modest in scale. The evolution of computer languages that allowed us to efficiently write code for machines was extraordinary and is responsible for the world that we live in today, including DeepMind and AlphaGo, etc.

So there are better ways of interacting with the world, and having a sharp edge is better than not having a sharp edge. I think where things get tricky, normatively, is when you start talking about refined cultural artifacts and objects. I know this is an interest of yours—different ways of reasoning, religiously reasoning about the world, or scientifically, or mathematically, or poetically, and so on. Are they like machine code versus Python? Is there a sense in which a certain culture has discovered a more efficient way of interacting with physical and cultural reality? I think it's a really interesting question, and I think that we know domains where the answer is yes. Having mathematics is better than not having it. There are certain things that we can do, like navigate, and put things on the moon, when we have it. So yes, it has incredible cultural implications.

Not many people think this way about the interaction between brain plasticity and the cultural accumulation of cognitive artifacts—especially in relation to collective intelligence and collective stupidity—which is rule systems that you've accumulated in the brain, which you thought you didn't need, and you didn't, but other people think you do, that oblige you

to interact with the world in a worse way than you did before. That happens a lot, as we both know.

So this is a brave new frontier. And I would be extremely interested in understanding it. In fact, one project at the institute that we just started is what we call the Law of the Legal Operating System of Society. Constitutions are a beautiful example of a memory system that encodes historical contingencies, events in the past, and our response to events in the past—hopefully with positive outcomes. We actually now have 590 legal operating systems, constitutions from around the world. We can ask, When do they work? When do they fail? What were their cultural implications? Which ones are more likely to lead to despotism? Which are less likely? I think this needs to be addressed, but I don't have the answers.

Harris: I wonder if there's a relationship between complexity and ethics or intellectual honesty? This just occurred to me: One difference between religious dogmatism and scientific curiosity is both the boundedness of the worldview that results and one's tolerance for ambiguity and complexity. For a dogmatist, the final answers are already given. Reality can't be more complex than what's spelled out in his favorite book. But for a scientist, or for just a curious person, the investigation of reality is open-ended. Who knows what we will learn in the future, and who knows how it may supersede or revise our current understanding?

When I think about the differences between cultures, I often notice what seems to me to be the most crystalline ones, which more or less tell you everything you need to know about the other differences between them. My favorite example of a culture that gets almost every important question wrong is the Taliban. I've been using them for years, but you could also think about ISIS or any society organized under strict *shari'ah*. I remember when my friend Christopher Hitchens described his reaction to the fatwa on his friend Salman Rushdie that came down in 1989 from Ayatollah Khomeini in Iran. His first reaction to this, when a journalist asked for his comment on it, was to say that it was a matter of everything he hated versus everything he loved: This single datum, that of a ruler of a state suborning the murder of a citizen of another country for writing a novel, encapsulated so much that was wrong with the culture.

I'm a father of two daughters. And when I think about the life I want to give them, and the kinds of things I delight in and worry about on their behalf, and when I compare this to the general attitude of men—and women, too—toward women and girls in traditional Muslim cultures, the Taliban being the ultimate instance, that difference betokens so many other differences. Take the most excruciating case: honor killing. With some regularity, a girl who gets raped, or who refuses to marry some old man she's never met, some second cousin her father picked out for her, or who wants to get an education, is killed by a male family member, who considers this a dishonor. I'm not talking about the behavior of a lone psychopath. I'm talking about someone who is psychologically normal in a culture that reinforces behavior that only a psychopath in our culture could possibly support.

This single difference, the treatment of women and girls, tells us almost everything we need to know about the likely differences on many other levels, intellectually and ethically, between that culture and our own. We know a lot about what a culture is *not* going to accomplish if it makes it a major priority to keep half its population illiterate and living in cloth bags.

Krakauer: I would say two things. The first is that the systems you are describing are intriguing instances of the persistence of rule systems, whose outcomes we would describe without hesitation as stupid—certainly in relation to the treatment of human beings. That is for me a genuine scientific problem. I know, as you do, that many people in those societies are deeply unhappy. These rule systems are imposed upon them. Why is it they're so persistent? By the way, in Western society, let's

be clear, women didn't have the vote until the early 20th century. But we realized the error of our ways.

The second is the implication I've already described, which is that rule systems leave an imprint on your reasoning in a very tangible form. If you are encoding a cultural form that is hateful or intolerant, just like the abacus, it is leaving an imprint on how you reason, on how you think about world.

What distinguishes a scientist from someone who has an orthodoxy? I guess it's enshrined in Richard Feynman's definition of a scientist as someone who believes in the ignorance of experts. That notion is the singular precondition for the possibility of science, which is a fundamental distrust in experts and expertise, including us. It has something to do with information, which as you pointed out, has something to do with uncertainty.

I've often thought that cultures tend to treat symptoms, not causes. You've described societies that barely are societies. My feeling is that we really should address these things somehow with a pedagogical schema that allows people to live with uncertainty—makes them happy about it, not unhappy about it. That reassurance should come in the form of possibility, not the lack of it. That's a deeper issue, and I think it's where our education of students is utterly failing. Because they're all symptomatically targeted, whereas what you're talking about—your notion of Hitchens's response, about all that I love and all that I hate—is this deeper issue of yes, we live in a void. The solar system is a dense bit of matter in an otherwise sparse universe. Do you delight in that? Or are you horrified by that? That kind of thing, that psychological profile, is what inclines you toward science or toward orthodoxies.

Harris: Expanding from there, how do you view the future of civilization or of our species in light of this basic uncertainty? Feel free to riff about various dystopian or utopian possibilities, but obviously, on one end there's the chance that we might destroy ourselves or that our global civilization might fail. There's also the possibility that we'll more or less engineer everything that's wrong with us out of existence and eventually export an unimaginably advanced culture to the rest of the galaxy. Most people seem to feel we're passing through some kind of bottleneck now and that this century is more crucial than most. Do you feel that way?

Krakauer: I do and I don't. You know, we've talked about this, and there are clearly characteristics of the 20th century that historically, with respect to our own species, were unprecedented. Population growth all happened in the past few decades, right? Computer technology, as we understand it, happened in the past few decades. Medicine that works according to scientific principles, as opposed to trial and error, is very new. Hygiene and an understanding of the implications of biological evolution in terms of the ethical treatment of each other and of nonhuman animals is new, and so on. So it's an incredible century, I think, in many ways. But in other ways it's not. You could argue that the first time we committed our internal representation to the world in the form of cuneiform lettering on clay tablets was a greater event in human history, with greater implications moving forward. There are times in history where extraordinary things have happened. It's hard to apportion differential weight to them.

Harris: I think you certainly can defend the claim that that was the breakthrough that enabled all the other ones we deem important. But what you don't have with the birth of writing is a technology that gives a single individual, to say nothing of a state, the power to destroy the species. I'm thinking of things like biological terrorism, or any other destructive technology

that can get away from us.

Krakauer: A lot of this is quantitative, not qualitative, right? Gunpowder was clearly extraordinarily important. Machine-guns as opposed to the cavalry, as we saw in the devastation of the First World War. It's a little bit like asking, Do we have excessive information processing now because we live in a computer age? And do we not see revolutionary transitions in human culture in the past because we think they can only be computational and atomic or biological weaponry in the present? But it's true that extraordinary things are happening. Not least, I think, the possibility in our lifetimes of the demise of the nation-state.

The kinds of social networks that are the prequels to territories and ultimately nations are different now, and the possibility of a true reconfiguration of terrestrial social systems is really intriguing. For many people who live on Facebook or in computer games, that has effectively already happened. It hasn't happened in the tax system. And it hasn't happened in terms of the electoral responsibilities. But it's happened in terms of how they live. So I do think there's a big change ahead of us. With respect to pessimism versus optimism, I believe in intelligence, and I believe in reason, and I believe in civilized discourse. I am frightened by unconditional optimism and unconditional pessimism. The two extremes have always upset me. And the extremes of politically correct and politically incorrect are both equally apparent, right?

So the middle ground has always seemed to people lukewarm and not inspiring. But that's exactly the bath I want to sit in. And somehow, moving forward, if we are aware of the distinctions, complementary, competitive, and the effects they have on the biological ability to reason, then we should be able to think about these devices as a community of civilized people and make decisions. One of my great fears, to be honest, has been what I see as a systematic erosion of human free will. And not free will as in where does it come from in a deterministic universe, but the moral implications of free will.

The example I often give is, free will is only as good as its empirical execution. That is, when you get a chance to exercise it. And it doesn't matter that you have it if you can't exercise it. So if ISIS came into power, it wouldn't matter if you had free will, because they would deny you the ability to exercise it. But we are voluntarily choosing not to exercise it. A few examples: "Netflix, what movie should I watch?" "Well, David, you watched these movies, so you should watch this one." "Thank you." "Amazon, what book should I read?" "Well, people just like you read books just like this." What this is doing, if you think about it geometrically, is contracting the volume of my free choice—under the economic pretense, in some sense, of allowing me to exercise greater free choice.

It is absolutely true that I could say no. But it gets harder and harder. Suppose we lived in a world where I wrote an app, let's call it Voter app. What you'd do is enter into this app your economic circumstances, where you live, your history of interest in politics, and it would tell you, better than you ever could, who you should vote for. And let's imagine the equivalent medical app—a sort of iWatch version 4; it would measure everything about your body that could be measured. And it would say, when you go to a restaurant, "No, you really shouldn't be eating an aubergine tonight. It's time for... whatever, a chicken sandwich or the reverse." I don't think that's alarmist. I think that over the course of the next decade, more and more decisions will be outsourced in this competitive form, such that what remains in our competence and in our hands will be a tiny particle of freedom.

Harris: I must say, I don't see those examples so much in terms of freedom. It's funny that you bring up free will, because listeners of this podcast will know that I spend a lot of time arguing that it's an incoherent idea. That's not to say that

everything else we care about is incoherent; obviously there are differences between voluntary and involuntary action, and not a lot changes when you get rid of the notion of free will, but a few things do change. My very last podcast had me debating Dan Dennett in a bar about free will.

Krakauer: It's important that it was in a bar, Sam. You had available to you mechanisms for increasing it.

Harris: Right. But, free will aside, I don't see those examples so much as a diminution in our freedom. There's certainly a siloing effect here. We're creating machinery that curates the available choices in such a way that it will reliably give us choices that we prefer to randomness, right?

Krakauer: Let me give you an obvious example. I'm a Western male, you're a Western male. You're probably wearing trousers and a shirt. The sartorial options available to you are extraordinarily small in comparison with world culture. And historically, the way we have chosen to adorn ourselves, in Persia or the Roman Empire or China, has been incredibly diverse and fascinating; yet now, as Western men, we all look like clones. I would claim that you're not exercising your judgment, you're being told precisely how to dress. And when you get to exercise your judgment, it is a very, very low-dimensional space in terms of texture and color that the manufacturers of clothing, purely for economic efficiency, have decided to give you. That's what I'm thinking about. You're absolutely right: You could do your own version of modern civil disobedience and say "No." But it's very hard for people.

What I'm concerned about is not inevitable, it's not deterministic. But unless we choose to assert our individuality and our constructive differences, we will, I think inevitably, become a clone species—not only in terms of the way we look and dress, but the way we reason. You asked me about my dystopian singularity—that's it. The optimistic future is the one where we say, "Enough. No more conformity, no more over-curation of what you think I should do and think." A kind of radical accession of diversity, a radical individuality that we somehow reconcile with a constructive communitarian drive. I don't think we've done that very well historically. How to be as different as we can be, but be congenial with one another. That is a positive future for me, but I think that's the path of great labor.

Harris: I think those examples are importantly different. I take your point on dress. It never occurs to me to even want to wear a kilt or something other than pants and shirt. In fact, I take that lack of imagination to an even greater extreme. I've decided to not even think about what I'm going to wear, so I basically have a uniform.

Krakauer: That's the omega point, Sam.

Harris: I'm the canary in the coal mine, sartorially speaking. But take the Netflix or Amazon recommendations. Normally, 20 years ago or so, you would walk into a video store into a bookstore, and you'd wander the aisles (leaving film and book reviews aside, but that's another curation process). You'd find specific covers or titles alluring—things would jump out at you, but it was largely a matter of happenstance. And there wasn't much information in the system to reliably promote any one among the thousands or tens of thousands of candidates for watching or reading.

So now we have things like Netflix and Amazon where, based on your reading or watching history, based on what millions of people very much like you have rated as enjoyable, you're getting various recommendations. I see a major liability here in getting ghettoized intellectually and ethically. This is happening online for most people who choose to follow others on

social media whom they already agree with. We get channelized, and the walls of the channel are getting higher and higher.

Krakauer: Sam, just to interject, let me be very clear: I am not a Pollyanna about the past. I'm saying something slightly different, which is that the tools we now possess, which are so incredible, should be allowing us to have freedoms that are unprecedented, not returning us to the ghettos of the past. So I'm with you. As you say, I used to choose my albums by their covers. That wasn't necessarily the most thoughtful thing to do. Although sometimes it works.

Harris: Spent a lot of time listening to Yes?

Krakauer: That isn't the kind of cover that works for me. But, yeah, exactly. So I'm not saying the past is great. I'm simply saying that if you develop a technology that could give you incredible freedom, why not use it to do that? I think what's so intriguing about the history of civilization and technologies is that with every new technology that offers some increment of possibility, it comes with the greater possibility of its own negation. So the bookstore example is a wonderful one. We were limited by our access to good bookstores—and most of them, quite frankly, were shitty, right? They had terrible taste. It was just endless shelves of self-help books that would help us better by burning to keep us warm. Amazon is a godsend with respect to access to books when you live in remote parts of the world. But what comes along with it is this all-seeing eye that wants to impose, out of largely economic considerations, constraints on what you do. And it's our job to maintain the freedom of the technology. That's all I'm saying. Let's fight the instinct of the technology to treat us as a nuisance in a machine-learning algorithm that would want to be able to predict us perfectly, and let's surprise it constantly. But yes, I have very little nostalgia for the past.

Harris: I'm now noticing the time, David. Let's open it up for a final consideration. How do you view the prospects for advanced intelligence elsewhere in the universe? Do you have an opinion about the Fermi paradox—you know, where is everybody?

Krakauer: I have an opinion. I don't think it's very well informed. I'm fascinated by space, and I should say that we're currently working on a new festival for New Mexico called Interplanetary, which is all about the future destiny of life on Earth and the universe. We could talk about that at great length. In terms of this calculation of life elsewhere, statistically, it's a real problem, because any well-informed statistical model has to have multiple independent instances for you to make an inference. The problem with our case is that there is only one. So you can't reason about this question statistically.

But you can reason about it in terms of physical law and evolutionary dynamics. Physical law, to the extent that we can measure it, is the same everywhere in the universe. And to the extent that biological mechanisms are emerging from physical law, there's nothing particularly special about Earth. By that kind of reasoning, based on mechanics, I think we have every reason to expect that life exists elsewhere. But you can't reason from the statistics, and that often leads to a rather fruitless discussion. Regardless, though, of whether or not there is life in the universe beyond our own planet, we have an intellectual obligation to populate it. That's where I stand on the matter. Why do we do what we do? I think that if I have any kind of quasi-mythical belief system, it's something to do with expanding the sphere of reason and sympathy into the world and beyond. If we could take the very best of what we've done and push it out into the universe, that would be an extraordinary thing.

Harris: That statement, that we have an ethical obligation to populate the galaxy is an interesting one, which I think will

strike many people as highly non-obvious. What are our ethical obligations to people who don't exist? It's interesting to consider. We know at a minimum that we have intelligent life on this planet that can enjoy a range of conscious states that can be beautiful and fulfilling. If we did something that canceled the future of the species—if anything is wrong, that is. I don't mean it's wrong in the sense that we'd be causing ourselves or anyone to suffer. We could kill ourselves painlessly in our sleep tonight, right? There would be no suffering, but we'd be foreclosing on potentially billions of years of happiness and creativity of a sort that we can't yet imagine. And that would be a terrible thing to do.

It would be great to have a technology or a device, something as simple as an abacus, that allows us to internalize a commitment to future generations. We're really bad at solving problems that have a time horizon longer than our own near future or our children's future at most—something like global climate change. We discount the pain of the future so steeply that we cannot prioritize a centuries-long problem at all, no matter how grave it is. If we could somehow make a commitment to the future more reflexive and more vivid, more emotionally and ethically salient to us, and internalize that—I think that's one thing we need. I don't know what it would look like, but the conjunction of your abacus talk and your saying it would be ethically problematic not to push forward into space in future generations made me think of this.

Krakauer: Well, I think it's a really intriguing and important point. I would claim that one reason so many of us are drawn to evolutionary thinking, to the ideas of Lyle and Darwin and Wallace, is that they do give us a sense of what time can do. For me it's extraordinary that over the course of billions of years, we've gone from a planet that looked like the surface of Mars, perhaps—it was lifeless—and now it's the Rolling Stones and Johann Sebastian Bach and Emily Dickinson. I think, as you say, that delicate, rare things should be preserved. And developing an awareness and a tangible ethics for that is really vital.

Harris: What's your view of changing the species in ways more radical than the mere happenstance of evolution? Genetically engineering changes that we presumably well understand into the germ line, or just allowing people to creatively change their genomes?

Krakauer: Well, first of all, I believe it's already happened. It happened with writing, and it happened in mathematics. I've already asserted that culture is a kind of collective inception event into the brain. We've been modifying ourselves forever, with nutrition or with exercise or in society. So the question is whether or not this will represent a radical discontinuity in the styles of intervention. I guess that comes down to the question of how much time it takes to change the system a lot. Let me be clear: We are going to modify ourselves, and if, for example, a pandemic emerged with a virus that had a morbidity rate of 80%, and someone had invented a modified CRISPR system to render you immune that required a change in the genome of each cell in your body, it would be adopted. Not only would it be adopted, it would probably be made obligatory. And that's not really that far-fetched.

So just in a matter of time, such things will happen, and some of them will be extraordinary. We'll probably be able to eliminate certain forms of cancer—not all. We will modify ourselves willingly and, I think, appropriately. The debate will persist, I guess, exactly the way it persists in the case of enhancement in sport—where to draw the line. And that comes down to a question of fairness, right? The ethics of fairness. So I do view the march of technology as kind of inevitable. But I would like to accompany it with reason. And one form of reasoning that's useful in these debates is to find precedence.

So when people talk about CRISPR, which is our currently the most powerful genetic engineering technology, it's worth

bearing in mind all the things we've done already to change genetics, either naturally or unnaturally, and what we've done to our microbiome in our diets and biochemically—which is a part of our genome, by the way, that we are utterly dependent upon. So it always helps, I think, to create continuity and reasoning, to find prior instances that we can use to think about the future. If at some point we were to colonize other planets, and those other planets had different masses such that the effect of gravity was greater or lower, they had slightly different compositions of gaseous molecules in their atmosphere, we would quite willingly reengineer ourselves. I actually think that's inevitable.

Harris: In closing, David, I'm going to ask you a question that you seem uniquely poised to find annoying, or even unanswerable, given what you said about IQ. But I've asked this of a few smart people on my podcast, and I think I'm going to demand an answer, no matter how much you recoil.

Krakauer: Demand it, man, demand it.

Harris: Who is your vote for the smartest person in human history? If you could put one human brain into the room to talk to the aliens, who would you nominate?

Krakauer: Probably John von Neumann.

Harris: That's actually quite an uncontroversial pick, given what I know about him. But do you want to say a bit about why?

Krakauer: Well, the thing about Von Neumann that's so incredible is that he created mathematical fields, physicals fields, computational fields, and social-scientific fields. That breadth and depth is almost unique to him. I wish I could pick several, but if I have to pick one, I pick him.

Harris: Just the stories about him, about the effect he had on the people around him, which included, arguably, more-famous and influential scientists and mathematicians—he was surrounded, as you know, by the most productive scientists of his generation. There are so many stories about the awe in which they held his ability to grasp and creatively interact with what they were doing, in real time, in a way that was just mesmerizing.

Krakauer: What was so incredible is that not only is there game theory, which he co-invented, and areas of quantum mechanics...

Harris: And the nuclear chain reaction stuff and meteorology....

Krakauer: Atomic physics...exactly. But as you point out, in addition to that kind of more traditional scholarly depth, he was frequently called upon, as you intimate, to solve problems that other people just couldn't even begin to think about. And he's such an interesting case, because he was this Jewish immigrant to the United States, left Hungary, worked on the Manhattan project, deep moral conscience. So he was a real 360, and I guess there aren't so many of them.

Harris: I don't know if you've heard this story about him, but on his deathbed, apparently, he was attended by the Secretary of Defense and members of all the branches of the military, on the off chance that he would say something useful about

nuclear deterrence. That's quite a testament to a mathematician.

Krakauer: Exactly, that's marvelous. When I was at the Institute for Advanced Study, I lived on Von Neumann Drive, which makes me very proud, you know.

Harris: Well, listen, David, it's been really a pleasure to talk to you. We could go on for hours and hours. If anything of interest happens in the world in the next few years, I will definitely invite you back to comment on it, because you're sitting at the confluence of so many interesting lines of inquiry. It's great to hear your thoughts on more or less everything.

Krakauer: Thank you so much, I really enjoyed it. I'd be more than happy to come back.

Harris: Before we go, David, I want people to know where they can find out more about you and about the Santa Fe Institute online. Would you direct them to websites and social media?

Krakauer: Absolutely. The website is www.santafe.edu, and I know we have Facebook pages and a Twitter feed. I should reassure all your listeners that we are completely reinventing our webpage, and by September, there will be something very beautiful to look at. In the meantime, we have to suffer through very dense materials. That's where you can learn about us and how you can engage with us.

Harris: And are you publicly funded? Or is it a matter of private donations to the institute? How does that work?

Krakauer: We are a not-for-profit. We're a 501(c)(3), and we're fiercely independent. We are funded, really, three ways: by grants, federal foundations, and restricted gifts. We're funded by something called the Applied Complexity Network. This is largely for-profit organizations—Google, eBay, Intel, Fidelity, etc. They've become affiliated with us because they're interested in complexity science in their own work.

Harris: Well, I encourage everyone to check out the Santa Fe Institute. And David, do you have a personal social media presence at all? Are you on Twitter, anything else?

Krakauer: I'm not. I'm so embarrassed.

Harris: That's why you're actually getting something done. Well, listen, once again, it's been a great pleasure to hear your voice, and to be continued...

Krakauer: Thank you, thank you.

The Waking Up podcast can now be supported on a per-episode basis on [Patreon.com](https://www.patreon.com).

Notes

Find this article online at: <https://www.samharris.org/blog/item/complexity-stupidity>