

AI and You

Transcript

Guest: Calum Chace

Episode 98

First Aired: Monday, May 2, 2022

Welcome to episode 98! My guest today is Calum Chace, a British futurist and keynote speaker about the impact of AI. Sounds familiar, no? We have a lot in common. Calum is the author of [Surviving AI: The promise and peril of artificial intelligence](#), and *The Economic Singularity: Artificial intelligence and the death of capitalism*. He is also the co-founder of a think tank focused on the future of jobs, the Economic Singularity Foundation. He called in from Spain to talk in this first part about what he means by the economic singularity, and how our socioeconomic system will cope with the changes that come from AI. And you'll also learn something about frogs. Let's get to the interview, with Calum Chace.

Calum Chace, welcome to *AI and You*.

Thanks, Peter, it's a pleasure to be here.

And so, give me the quick, brief description of what it is that you inhabit the space of as a thought leader?

I describe myself as a speaker and writer about the impact of AI on all of us.

And I have here, one of your books, *The Economic Singularity*, which I very much want to get into. First, let's visit some of your history and this because I'm always fascinated by how people came to the field of AI. Very few people start out there unless they're relatively young and you started out with, I think, the BBC and the *Financial Times*. Tell us about where you did start and how it influenced what you're doing now.

So, you're quite right, my first job, my first sort of career job was as a trainee journalist with the BBC and later in life, I had a column for the FT. So, journalism and writing has been a theme in my career, but most of my early career in my first 30 years, I was in business, first as a marketer, and then 15 years or so as a strategy consultant and then I ended up running a series of small companies. And then I retired, and got to do something that really interested me, which was reading and writing about AI. My introduction to AI - as you say, I'm not an AI researcher, something that most people do that a lot younger - my introduction to AI was, firstly, through science fiction, which I read as a kid and there's lots of lots of science fiction stories about AI. And then like a lot of people, Ray Kurzweil opened my eyes to some surprising possibilities. So as a kid, I always thought that there would be AIs as smart as humans. But I always thought it would be thousands of years into the future and there was no chance I'd ever see it. Ray Kurzweil opened my eyes to the power of exponential change and the exponential growth of our technologies on and the remarkable possibility that we might get super intelligence, say in this century. He thinks that that will happen in 2029, because he's been saying that for a very long time. I think that's probably overly optimistic, but he's a smart guy. So, when I came across this

idea, which was in 1999, I read his book, *Are We Spiritual Machines?* So, I thought this was a startling idea, and something that many more people should be aware of. So, I wrote a novel about it, because I thought it'd be a good idea to have a movie on the subject and I thought, I can't make a movie, but I could write a novel which could become a movie; it didn't. I discovered I'm a better nonfiction writer than a fiction writer. So I turned the novel into a nonfiction book and that's how I got into the space.

And you mentioned Kurzweil, of course, he introduced the singularity, the term *singularity*, which your book title is echoing, and is also very vocal about longevity, which we'll get onto later. But Kurzweil's singularity, and he copied the term from cosmology. So he said, there will be a singularity - this is one of those things that just makes me envious in the thought leader space of anyone that can grab that chunk of a meme and run with it - that there will be a singularity in 2045 and what that means is that the pace of development will at that point become effectively infinite, driven by artificial super intelligence and other exponential technologies. That to our perception, everything will be happening at once. Now, the term *singularity* in cosmology refers to the thing at the center of a black hole, which is where the laws of physics break down, you're dividing by zero. So it has a very sharply defined boundary; infinitely sharp, and so to refer to either Kurzweil's singularity or your singularity invites this idea that it is similarly sharply defined; and I've always wanted to ask Kurzweil: So what would we notice on the morning of the singularity, like two hours before, the night before, or even perhaps less frivolously, would we be able to tell when it was going to happen two years earlier because... might like to make some sort of preparation. Define the economic singularity for us and distinguish it from Kurzweil's and then tell us, does it also have a sharp edge?

Yeah, it's a good question. So, the word "singularity" has got all sorts of associations and I've wrestled with the idea, with the question of whether it's a good idea to use it or not. It's borrowed from maths and physics and as you say, in maths and physics, it's a point where a variable becomes infinite and the classic example is at the center of the black hole where the gravitational field becomes infinite and what happens then is, as you say, the laws of physics break down. That's the classic maths and physics definition of "singularity." That's not really what Kurzweil means by it; he really means a point at which everything changes, that a total break with the past. Now, somebody counted up a number of different ways that people have used the term "singularity" in the, broadly speaking, future space, and got to the number of 17. My use of word singularity simply means it's the biggest change that you can have. So, it's bigger than a transformation, it's bigger than a revolution, it's a really complete break. And the arrival of super intelligence will be such a complete break. The life of humans ever after will be completely different from the life of humans before. And I borrowed the term for the economic singularity because I think point when we have to admit that most humans will never have jobs again, because machines can do almost everything that we can do for money cheaper, better and faster, that is going to be a total break with the past. I think it can be a very good thing; I think a future world in which humans aren't able to do jobs and can't do jobs should be a world in which we do whatever we want to do instead. That's going to require some big changes in the way we run our economies, obviously, it's going to require significant transfers of resources from the

people who've got most of the resources to the rest of us, but I think that should actually be - not easily achievable, but definitely achievable - and we should have an economy of abundance. So, I use the term "singularity" simply to mean the biggest possible kind of change, which is what I think Kurzweil does in the context of artificial intelligence. His singularity is a bit odd, you're right, he says that it will happen in 2045. I've never been quite sure what he thinks is going to happen between 2029, when he thinks we'll create artificial general intelligence, which is an AI with all the cognitive abilities of an adult human; the 16 years between then and 2045, I've never been quite sure what he thinks is going to happen in that period. I've read all his books, it's not really very clear but for some reason that 2045 is when he thinks that if you like the singularity, that this super intelligence will transform our lives completely. So, there is definitely going to be a singularity, I think, when we get super intelligence; personally, I think it will happen later in the century and I think before then we'll get an economic singularity, which is what I wrote about in that book, which is joblessness, and the economy of abundance and that completely different lifestyle we'll all have. And I've recently been persuaded - a friend of mine has been badgering me for a while - so, there aren't just two singularities, there're more, and I think he's right, actually, I think there's going to be a longevity singularity, and I think there's going to be an agricultural singularity, we could get to those if you want, but as I say, I use the term "singularity" simply to mean biggest kind of transition, biggest kind of change that you can have.

And - I could be wrong on this - I think Kurzweil drew this distinction that in 2029, he expects or predicts *human level* artificial intelligence, and not necessarily artificial *general* intelligence and there may be a distinction between those. Certainly, when I was talking with Stuart Russell, we were saying how we lack a vocabulary or a science to draw distinctions about artificial general intelligence, which are going to become necessary because we increasingly are doing things that eat part of that space, but not as much as we would like or hope. But I want to say that you say that Kurzweil's singularity is more important, but yours is more urgent. If I got that right, could you speak to that?

Yeah, that's absolutely right. The arrival of super intelligence will be, I believe, the most important event in the whole of human history, because once we have a super intelligence on this planet, we will then be in the position that chimpanzees are in today. Our future will be entirely dependent on the decisions of that super intelligence. We won't have a say, in our own future. And personally, I think our only safe future at that point is to merge with the super intelligence, whether that be by uploading our minds or some other sort of brain computer interface, which I think is broadly speaking, what Kurzweil thinks will happen as well. So, it's that you can't really get more existential than that. The downside of the creation of super intelligence - and a lot of people poopoo this, but I think it's very real - if you have a super intelligence - might not even be a conscious one might not be a malicious one, but if its interests are not aligned, if its goals are not aligned with our interests, it could easily be an existential threat to us. So, I think the arrival of super intelligence is the most important singularity. The economic singularity, I believe will happen sooner. If it doesn't happen sooner, then it won't happen at all, because it'll get wrapped up in the, in the technological singularity. And therefore, because it is coming sooner, it's more urgent to deal with it sooner. And I think that if we get the economic singularity wrong, if we don't get the economy of abundance, which enables us all to live a life doing whatever we want

to after joblessness arrives, then we will probably get massive social breakdown. And it may not be an existential threat, but it would certainly very likely wreck developed societies. So, it's urgent, it's very, very serious; it may not be existential in the way that the technological singularity is.

I want to dig into this because of this is, as you say, urgent, and massive social breakdown seems to be something we're increasingly flirting with these days, starting from the effect of disinformation campaigns on national votes, and now obviously, with the war in Ukraine. And, I'll just go straight to the issue that is most central in this. Artificial intelligence and other exponential technologies create dividends that anyone predicting their effect decades before would say this will result in a cornucopia and everyone will live happily ever after that, what we find when we get there is that that's only true for the people who invent and hold those technologies, and they don't trickle down to the other people. There seems to be this - that capital attracts more capital, and so we see the ratio of CEO to worker pay increasing, and the concentration of wealth and the upper point fractional percentages increasing, and only begrudgingly going to the rest, the lower parts - and I'll give you opportunity to rebut wherever you disagree with this, because I would definitely want to get into that - but the fundamental question I'm aiming for is that when artificial intelligence creates a result that is a financial dividend, it goes into someone's pocket and in order for that to be shared, among a lot of other people, you have to take it out of their pocket, which would involve some kind of different paradigm, it seems, from the capitalist one that we have now. That's a lot. Okay. Take as long as you like.

Okay, well, yeah, you're right, I do disagree with that. If you look at the lives of people today, even poor people in the developed world, they are better lives than the lives of kings and queens 300 or 400 years ago. Kings or queens 300 or 400 years ago, died younger, died, less healthy, less well, were massively less well informed, massively less comfortable. couldn't travel as much. On every metric, a poor person's life today in the developed world is better than a monarch 500 years ago, maybe not, you know, obviously not in terms of status and a lot depends on how you compare yourselves. Emotionally, a lot of our wellbeing depends on how we compare ourselves with our peers and so on. So monarchs beat us on that score, but in every other objective score our lives are better than theirs. And there's two reasons for that. One is technology and the other is capitalism. I have an unapologetic defender of both. I'm a techno-optimist. It seems to me really obvious that technology is what has made us the preeminent species on the planet and has improved our lives enormously. As we've developed it, and capitalism, the industrial revolution of capitalism is what's given us this incredible improvement in the quality of our lives. And it really is that recent, I mean, if you look at world GDP, from the time of Christ, say, it was pretty flat and then, when the industrial revolution came along, it rocketed up, it's an exponential curve. So, I don't buy the argument that capitalism means that the benefits of technology accrue only to the wealthy and it has to be taken from them. It's a fundamental fact, I think, that you make more money as a vendor of any kind of technology, by selling billions of units at a low price than you do by selling a few units, encrusted in diamonds to the uber wealthy and it doesn't take long. I don't really like the word "trickle down," because

it has Reaganite, Thatcherite connotations which makes me uncomfortable and [that] economics I think is pretty much rebutted. But it isn't true that technology stays in the hands of the wealthy, if you go around the world, to the less developed world, as well as the developed world. Anybody who's in the middle class, and that's most people these days, all around the world has a smartphone. You know, you'll see people glued to their smartphones in train stations in Nairobi, just as much as in London and New York, and Vancouver. And as a friend of mine says, you know, we talk about the "digital divide," well, there isn't a "fridge divide" in America, for instance, there isn't an "air conditioning divide." Most people have got air conditioning, I was listened to a podcast this morning saying that the only people in America don't have air conditioning is prisoners, which is a terrible state of affairs. But technology does permeate the whole of society and the benefits of AI will permeate old society. Capitalism has flaws, for sure. It needs reform, it needs constant reform. Businessmen and women are always trying to cheat. They're always trying to form monopolies and cartels. That's in the nature of things and regulators have to stop that happening. But it is an amazingly powerful way to allocate resources and to encourage innovation.

Well, then, perhaps I'll change the tone of the conversation from *alternative* to capitalism to *improvement* to capitalism. What do you think about the Gini Coefficient of the difference, the ratio between the most well-off and the least well-off in some region? There's evidence that inversely correlates with the happiness of people, because the less well-off ones, relatively speaking, are always comparing themselves to the most well off ones and if you're living in a tent, and see someone's super yacht go by on the river, then you feel that more acutely than a smaller difference in a society that's done more to elevate people out of homelessness, for instance. Should we evaluate our success in the dissemination of the benefits of technology by the relative difference between individuals in a society?

Yes, I think a hallmark of a civilized society is how well it treats the least well-off people in its community, and one of the things that needs to be done to capitalism, because it doesn't do this on its own, is to create a safety net for people who are unfortunate enough to fail to prosper; there will always be those people. The Gini coefficient is a rough and ready measure of inequality; some people don't like it, it seems to me to be a pretty sensible measure. And in Europe, where I'm based, it hasn't really changed in the last 30 or so years. Wobbled around a bit, hasn't really changed. In America, it's got a bit worse. And I understand that really most of the driver of that is the uber wealthy, the CEOs and the founders of businesses. Now, I don't really think - this is probably an unpopular view - that people feel better or worse off, because they compare themselves to Bill Gates: I certainly don't. I compare to myself to people that I grew up with, and people from similar backgrounds to me. And I try not to do that very much, I'm reasonably good at not doing that too much. So, I don't worry about too much. But I think people do worry about keeping up with the Joneses rather than with the Gateses. And so it doesn't bother me really, that Elon Musk is worth whatever he's worth now - 200 billion, I can't remember. He's kind of on a different planet for me and I don't really mind that he's got that sort of level of wealth. He lives a lifestyle that I couldn't possibly aspire to; it's not really in my range. Where the Gini Coefficient is bad, it tends to be in developing societies. And China,

which is currently trying to solve its massive inequality, its really sharp inequality, making a serious attempt to restrain that. India's got really bad, really extreme inequality, Brazil's got really extreme inequality; that's where it tends to be. And it seems to be a pattern that countries go through. As they develop, as they industrialize, inequality gets worse, pollution gets worse and then increasingly, you get a middle class developing, and the real extremes of inequality start to improve, and pollution starts to be cleared up. So it seems to be something that there's a development path that countries go on. So, I take the often-unpopular view and minority view that capitalism is a good thing; inequality isn't as bad as we think it is. One scale in which inequality is certainly not as bad as it was, is globally. 20, 30 years ago, it was the case that most of the world was very, very poor and a smallish part of it in Europe and America, mostly, was relatively very rich, and a huge middle class has developed, obviously, in China, obviously, in India; less obviously, but also true in Africa and in South America. So the world is globally much more equal than it was 25-30 years ago, because countries are developing.

Come the economic singularity, what will the flow of capital from technology dividends be doing? What are the risks and benefits of what will happen at that time?

If we play our cards right, what should happen - and this is kind of already happening is goods and services will get cheaper and cheaper and cheaper. It's often observed that goods like cars and televisions, chairs, are much, much cheaper, relative to an average person's income than they were 20 or 30 years ago. And as somebody who grew up in the 70s, I can tell you, that's definitely true. No, when I grew up in the 70s, compared to the way people live today, we were all poor and not just poor, things weren't very good. Cars leaked oil, and they broke down all the time. If you drove any distance on a motorway in the 70s, you'd see loads of cars broken down, people with the hoods up trying to fix them. Television was black and white and hadn't long been color and they were in the UK, they were two channels, then there were three, and they were rubbish and they weren't on during the day, and they close quite early at night. And the quality of programming was appalling. And now you have this cornucopia of information and entertainment, you have consumer goods, which are massively better. So, we're on a trend to making goods and services, quality per pound cheaper and cheaper. There are some exceptions, healthcare education, particularly are not getting cheaper. And that's a problem that we need to solve, and I think AI have a role in that. And I think that AI will accelerate this process and will get to a point where the cost of all the goods and services that you need for a very good standard of life will be almost free. Not completely free, but almost free. And that means that when we get to the point where machines do almost all the jobs, because they can do them cheaper, better, and faster than we can, it won't cost very much to make all the goods and services that you need for a very good standard of living available to everybody. And that's the economy of abundance. That's what I think will happen if we play our cards right at the time of the economic singularity. The big threat to that, the big risk, is that people see technological unemployment coming, and they don't believe in the possibility of an economy of abundance. Right now, there's a very small number of people who believe that we could have an economy of abundance. I know a few people who do believe that, but there's not many of us, and right now, if you persuaded 70% of the world's population that technological unemployment is definitely coming, you'd have a really big panic. And what happens when you have a really big panic, you get bad politicians

getting elected, you make bad decisions, and eventually it ends in war. That's where it always goes. That's where populism always goes, if it isn't overturned. So the big threat to the economic singularity, I think, is panic and I think we need more people to understand the immense power of AI the immense power of exponential improvement in technology, and the wonderful possibility of an economy of abundance.

I also grew up in Britain in the 70's and I think I would attribute the breakdowns on the side of the motorway to British Leyland more than any other factor, because of their reputation. To be perhaps simplistic about it, but to couch it in brief terms: if I am technologically unemployed in the future, even though the things that I want then will be cheap, they won't be free; will I be able to afford them?

Yes, because there will have to be a transfer from whoever owns the assets and is receiving income to everybody else. But if we do achieve the economy of abundance, the cost of that transfer needn't be onerous and it's just a normal taxation. So I envisage a world - I mean, there's different ways this could pan out, with different models. My preferred model is that we don't get rid of the market economy. The market economy is a really efficient way of distributing, allocating resources. And so I imagine a world in which there are still people making money and some of them are making fortunes. There will be Elon Musks in the future. Most people won't be engaged in that because they don't need to. I'm not sure I would be engaged in it, I would get on with studying AI and looking at what's going to happen in the future and looking at history and traveling and the things that I enjoy doing. And so people who are still engaged in a market economy will be making fortunes and won't have to be taxed very much for everybody else to receive enough resources to have a really good standard of living.

Do you think that we need to change the method of taxation? Just to make it clear here, I'm in broad agreement with you there; not an economist, so, I'm looking to you to learn from you about this. Bill Gates proposed a robot tax and I was never able to make any sense out of that to the point where I wasn't even sure that he was serious. And will the way that we currently do that through capital gains tax, sales tax, value added tax; will that will continue to work?

I think it probably will. I don't see why we need to have a total conceptual rethink of taxation. And I agree with you, by the way, I think he was basically being flippant, I think he was simply saying, look, if we've got a world in which AI is generating, virtually infinite wealth, surely, we can redistribute enough of it to make everybody comfortable. It's a terrible idea to say we'll just tax robots, because if you just tax technology, then you lose innovation. It just stops. So yeah, I don't think he was being serious.

Plus, it's really hard to define what the robot is, or where you do that. You know, Hollywood makes a lot of our thinking or influences a lot of our thinking about artificial intelligence, of course, and its persistent meme is that artificial general intelligence arrives in a thunderbolt when some mad genius invents it in the lab, and things change overnight. We're seeing that's not true, and that we are arriving at what we would have thought of artificial general intelligence by dribs and drabs with things like being able to play Chess and Go and now we

have GPT -3 doing things that are creatively within certain parts of our creative expression. And does that gradual arrival of artificial general intelligence, is that a good thing or is it a bad thing in that it creeps up on us and we start to be like the frogs in the tub of water that's coming to a boil?

Yeah. Interesting facts about frogs in water is it's not true that they sit there and get boiled, and this has it been known for a long time but it's a good metaphor. And it's one of the interesting effects of the internet actually. Before the web came along, there was a book called, I think it might have been called the *Age of Unreason*, written by a very famous business writer whose name for the moment escapes me, Peter Handy, I think he was, and he had the image of the frog boiling in water as the motif. And back then it was quite hard to check whether frogs do sit there and just get boiled. And then the web came along, and it became really easy to check and as it turns out, frogs don't. So scientists have known that for a long time, because back in the 18th century, people tried boiling frogs in water, and guess what, they jump out. So I think it's probably better for us if AGI arrives gradually because if it arrives suddenly, and - just kind of take one scenario, one day AGI what was one day the most advanced AI is very capable, but it's not conscious, it's not making up its own goals. Then the next day, it wakes up, starts making up its own goals, has a look at us and thinks, "you might possibly be a threat to me, and I don't really like you." That's a really bad outcome and so if that happens, suddenly not good. If you can see signs of it happening over a few months, you could take steps to stop it. Maybe, although that wouldn't be easy, but you possibly could. So, it might well be that it arrives gradually. But you know, we really have no idea what AGI will be like. Its mentality may be totally different from ours. Its way of thinking, if and when it gets to the point where what it does its calculations can be described as thinking, may be completely different to us, may be alien to us. How it arrives, what it's like when it arrives, is still totally unknown to us. But I do think we should be thinking about it so that we can try to make sure the outcome is a good one.

That's the end of the first part of the interview, part 2 will be next week. Obviously there are some huge questions around how capital will flow when technology like AI is creating far more of it and you've heard some deep thinking around that on this show, but it's obvious that so much more remains to be done.

I've got a letter from listener Paul, who writes:

I imagine people in the future may have humanoid robots as companions or servants and we will be able to order them with human behavior characteristics dialed to any level we wish. You've no doubt heard of some robots or chatbots going " bad ". I seem to recall an example of one robot that was asked particular questions causing it to go toward a psychopathic or sociopathic trend. Will robots of the future need to be calibrated from time to time to make sure the desired behaviors are still within the range that they were originally programmed at? And will these calibrations need to be audited by governing bodies?

Great questions, Paul. The chatbot I think you're referring to is Tay, from Microsoft, who deployed it around 2016. I believe it stood for Thinking About You. It was designed to learn from the conversations people had with it. Now it was based on another bot called Xiaoice, which had been successfully trained in China, so they thought this one would work just as well, so they put it on Twitter. Well, let's just say

that things did not go well. Twitter in the USA is a very different beast from whatever they used in China. Before Microsoft pulled the plug on it after only 16 hours, users had it trained to spew racist and sexist insults and repeat Nazi propaganda. Kind of reminds me of the experience of Hitchbot, which was a robot that wasn't able to move, it could just tell people that it was traveling across Canada and would they give it a ride? And it was successfully transported across Canada within the space of some weeks. I believe people would take it some distance and then leave it on the side of the road again where it had a thumb stuck out. Then the makers decided to try it in the US. It didn't make it out of Philadelphia before it was dismembered. The head was never found.

So it's fun to read all kinds of social commentary into that, but at least one of the points is obviously that different environments have different rules and different effects. Now, Paul's point about calibration is rather interesting, because a lot of the narrative about the robots of tomorrow implicitly assumes that they will be self-governing after some point, but Paul points out that calibration – and re-calibration, sort of like getting a front-end alignment on your car – is an obvious mode of operation. Microsoft didn't attempt any sort of recalibration of Tay, they just buried that thing as fast as they could. And it's an even more interesting point about how the governance of that calibration would be done. Some sort of third-party standards body is exactly the way any question of the safety of a machine in any proximity to humans has been managed. Think about everything from microwaves to airplanes or chainsaws to metal presses. There are standards bodies like the BSI, CSA, UL, and the FAA to name a very few, that have inserted themselves into those deployments.

But who would do that for robots, or chatbots? No idea. Would they emerge as some sort of government-sponsored organization, or a spinoff from academia? Still don't know. It's clear that we're pretty early in this cycle. If you look at an environment that's been going a bit longer than interactive AI – social media – it's clear that there are dynamics like disinformation that warrant some sort of regulation and yet the industry is still entirely self-governing and there's no hint of a body that would certify, say, some part of the metaverse as meeting standards for the truthfulness or harmlessness of the information you'd encounter there.

So, Paul, great idea, and since no one seems to have done it yet, maybe you could start the ball rolling and let us know how it goes?

In today's news ripped from the headlines about AI – wow, this one's got some big numbers. Kind of our benchmark for comparing computers to the human brain is in the number of synapses, or connections between the either 86 billion or a hundred billion neurons in the human brain depending on who you listen to. Each of those averages a thousand connections which means you're talking about a hundred trillion, in round numbers, parameters that are equivalent, or so we think, to the weights in an artificial neural network that go into computing the activation function on each node. Or, the number of parameters in a transformer like GPT-3, which has 175 billion, which makes it about one five-hundredth of a brain, and the biggest transformers since then are around a trillion parameters. Well, now the UK company Graphcore, which designs AI chips, announced plans for a 500 trillion parameter computer, which they intend to mass produce for \$120 million each. Which is cheap compared to the cost of today's supercomputers. But that would put it at five times one measurement of the human brain. Does that mean it will pass the Turing Test and run for office? No – remember Stuart Russell saying on this show in episode 86 that we're currently software-bound on artificial general intelligence anyway, so what we need is the software. Graphcore calls it the Good Computer, after Irving Good, Turing's

assistant who made some significant statements about our future with evolving AI. They say it'll be ready in 2024. Can't wait to see what happens with that.

Next week we'll conclude the interview with Calum Chace, when we'll talk about AI in the metaverse and how that virtual reality might evolve to become more useful and impactful. That's next week on *AI and You*. Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

<http://aiandyou.net>