

Anthony Comegna ([00:24](#)):

The singularity is near, or at least so we're told. Within the next 30 years or so, our world, our society, the very nature of our lives may change so sharply and so significantly that the past before the change is no longer recognizable. Looking out over the astonishing and terrifying changes that seem so imminent in the near future, it is easy to sometimes get lost in the doom and gloom and forget that we, classical liberals, will also be part of building that future. So what will we make of it? Well, Economist James Caton joins us to sort it out. All right, Jim, thank you so much for joining us here. And this weekend you'll be leading a discussion for us, a Graduate Student Discussion Colloquium on Artificial Intelligence and Liberal Futurism. And now that's sort of a new topic from what IHS normally does. And I'm wondering then, what in particular got you interested in this subject?

James Caton ([01:38](#)):

Well, it's good to be here and it was sort of a roundabout way of getting to even being familiar with AI. Certainly when I started my graduate work, I was not expecting to go and research AI. That wasn't my intention. When I arrived, I was at George Mason and my intention was to study Austrian economics and monetary economics. And my hope was, after having, I earned my Master's Degree at San Jose State and I studied under the likes of Jeff Hummel, Edward Lopez, among other talented individuals over there. And having learned quite a bit of economics beforehand and having been familiar with some critiques of the neoclassical system. I was anticipating finding a well worked out, fully worked out, I should say, which probably was too high of an expectation for any system, social system, economic theory, which certainly there's lots of gems and lots of treasure there, but what I ended up finding was that I had to start looking at what rationality was.

James Caton ([02:51](#)):

We have a lot of discussion about rationality in economics, rationalities in terms of prices, efficient markets, economic decision making. You might hear the word all the time, oh well, people aren't really rational. But I'll venture a guess that most of the time when you hear that, the person saying it doesn't have a definition of rational, which was precisely the problem that I was confronting when I started really digging into my graduate work toward the end of my first year at George Mason, was how do we actually truly represent economic decision making? And what you might think of ultimately is entrepreneurship that drives the market process. What's the appropriate theory for that? Everybody seems to be working on a lot of different margins. There's a lot of good research there, but we're still generating consensus. And so my gateway into AI, that is to say, was through the gateway of rationality.

James Caton ([03:51](#)):

Which ultimately, artificial intelligence systems are trying to construct rational systems that can interpret the world and make predictions about it.

Anthony Comegna ([04:01](#)):

So now, I take it that you use sort of, let's say, a classic Misesian definition of rationality. It's the ability to apply means to the fulfillment events.

James Caton ([04:11](#)):

So this was the most fascinating thing when I started looking at the artificial intelligence research. So, for those who don't know, my background is working with agent-based computation largely, and tying that to a well-rooted theory, economic theory, economic theory with roots in rationality. Okay. That was

where I ended up. And Mises offers this definition of rationality, of agent action, rational agent action, as being simply value driven. Any action you take is action that you expect to generate more value than if you didn't act, right. And so when you make a decision, this is just classic opportunity cost, you have the world with your intervention and you have the world without your intervention. If you think of action as intervention in the world. And you are, your action is either going to shape the world to be one thing that you would expect given your intervention or to allow it to be what it otherwise would have been.

James Caton ([05:16](#)):

Now, of course there's all sorts of imagined opportunities, but at the end of the day you choose the highest valued one in terms of if you choose to explicitly intervene, right? This is basically what rational systems are doing. Trying to order the world to use logic to make a decision about it. Sometimes that decision is as simple as just rating the data or predicting the data, coming up with a model to predict the data. But with humans it's looking at all the information, gathering it and trying to make the choice that's going to generate you the most value.

James Caton ([05:52](#)):

Artificial intelligence follows this schema, so I had read all of this Mises, all this Austrian economics and I started reading about artificial intelligence and the whole idea was that the systems interpret data or interpret an environment. They rank the valuation, they value potential actions and then the system takes the action that is considered to be most valuable by the rating. This sounds a lot like human actions, kind of a mix between your neoclassical utility theory. But it doesn't have to follow precisely those precepts.

James Caton ([06:29](#)):

AI systems are very diverse. And this is one thing this weekend, as we get into discussion I expect, and I hope people are saying, are discussing, is that there are many, many, many different kinds of artificial intelligence systems and they're not totally separate from one another either. They can be used in tandem.

Anthony Comegna ([06:50](#)):

Well, yeah, that brings me to my next question that popped up here, which was, so what are the differences between rationality and intelligence?

James Caton ([07:01](#)):

That's a good question. So rationality is certainly aimed at evaluation. I think rationality is a bigger claim than intelligence, and there's definitely a lot of overlap. But rationality, as far as I have come to understand rationality is of evaluation, right? You have a logical system that coherently describes the world as in, it's either not generating conflicting evaluations or if it is, it has a means of overcoming the conflicting evaluations to actually make a decision. And then rating the world. It's able to orient the acting agent who is supposed to be rational in terms of actions and expected future outcomes. Okay.

James Caton ([07:53](#)):

Now, maybe that's not the only definition of rationality that you can make. But in terms of if you're asking me if an agent is rational, that is how I would be describing their rationality. Intelligence. Well,

that's more like just the framing, rationality uses intelligence, right? But that's just to say that in order to be rational, you need a model of the world that's going to interpret that world.

Anthony Comegna ([08:21](#)):

Yeah. So it sounds like maybe intelligence is something like the ability to amass information and then rationality is the ability to use that information towards some productive end.

James Caton ([08:33](#)):

Yeah, I think that's a nice summing of what we have here. Now of course, right, you may be get into this discussion of intelligence and wisdom, right? And this is definitely, I think, at the forefront [inaudible 00:08:42]. Well, we have some really intelligent artificial intelligence, but the real struggle is to get some really wise artificial intelligence. And that is a struggle, especially when you start thinking about more than just a narrow domain where it's already difficult enough.

Anthony Comegna ([09:02](#)):

Now, before we get too far in the weeds here, let's back up and talk about a big conceptual thing that I'm sure we'll encounter right away this weekend, which is the singularity, this idea of a technological singularity. Go ahead and introduce that concept to our listeners.

James Caton ([09:20](#)):

So, the singularity, it's something of a fuzzy concept. It has some definitions, but doesn't only have, as far as I understand, it doesn't have only one definition, but you can think of it generally. Some will describe it as a unification of human and artificial intelligence or human intelligence and technology. I think that's one simple way to think about it. So yesterday in the news, Elon Musk is saying that there's a chip on the way to implant in your brain and to increase the efficiency of your thought processes. Which sounds, one, a little bit frightening, and two, forgive me, but I'm somewhat naive. So that's one way of thinking about the singularity.

James Caton ([10:00](#)):

And other is just to think about the rate of progress of artificial intelligence and technology. So this was originally John von Neumann in his description was thinking about something like a takeoff point of technology, where technology is developing so quickly and that means knowledge is developing so quickly, that it's sort of, it's throwing into disequilibrium what we think of as human life, right? The way of doing things is no longer keeping pace with the advancement in technology.

Anthony Comegna ([10:34](#)):

Yeah. And the idea is that we're right now on the cusp of, or perhaps in the middle of, an artificial intelligence explosion that creates a new period of singularity. And the last one would have been the industrial revolution. 1860s, '70s, '80s, and so the upshot of that is something that most people are probably familiar with. Deirdre McCloskey's hockey stick model, this incredible increase in production that just transforms life for people and still is in that process. And so, I wonder though, some people, like Ray Kurzweil has put his predictions at, say, 2045 is when we're going to have this intelligence explosion. And like you said, there's no single definition out there of this singularity. Some people take it to a religious level. I've heard it called the rapture of the nerds.

Anthony Comegna ([11:34](#)):

And so my question to you though is have you heard a definition of the singularity that is particularly compelling to you? And do you think that any of the timelines that you've heard suggested for it, many of which came up in our readings here, do you think any of them are legitimate?

James Caton ([11:55](#)):

Well, I like the idea of something like, not that I like the idea of general artificial intelligence, not to a particular task, but generally. So if we're talking about how, or we're talking about the AI from I, Robot kind of AI, where we have this intelligent system that seems like it has sentience. It's not just inputs and outputs, it's actually having a conversation with you and seems to have desired ends, seems to maybe even be able to choose ends or those ends can transform, right? It can interpret many different domains. Right? This is the idea of general artificial intelligence, very Hollywood kind of a presentation of artificial intelligence. But something that's been talked about alongside of this idea of singularity and throughout since the beginning of the generation of artificial intelligence, which you can frame roughly in the '40s and the '50s and the Turing Test, right?

James Caton ([13:00](#)):

I think this is a good starting point for thinking about, well, how much closer are we getting to something like that? Maybe the Turing Test is a signpost on the way to singularity, if there is a singularity, and we can remain skeptical about whether ... how we define singularity, whether there actually is one. But if we have something like a general artificial intelligence, we could start asking the questions of, who's serving who? Or who's serving whom? Right? The robots, the humans serving the robots, the robots serving the humans? Is it not really quite clear altogether? Are we serving each other now? I think ... so, you mentioned Kurzweil and he has this idea about something like AI emancipation or AI self-ownership, right. So do we get to a point where robots are no longer capital owned by humans but rather they own themselves, control themselves and are subject to all of the same sort of regulations that humans are generally held to rather than responsibility being taken by the owner of the AI?

Anthony Comegna ([14:11](#)):

Yeah. That's another question I had for you here and that I have for everybody coming up this weekend is does this possibility of artificial intelligences, AIs, individual intelligences walking around embodied in different devices perhaps or just in the cloud out there as code or something that I don't quite understand. Would their existence throw some serious monkey wrenches into certain cornerstones of classical liberal thinking? Like, for example, private property and I mean, you just mentioned this, this idea that a transition point would be reached where these things, these robots or computers are now sent seemingly sentient. They're no different essentially than you or I. It's just that their casing is a little bit different than ours. Ours are made of meat and theirs are made of metal. But in other important respects, they're the same as us. So do we have to then stop treating them as property and how strongly are people likely to resist a change like that?

James Caton ([15:26](#)):

So I'm actually quite skeptical about predictions to that point about the singularity. And I think there are some major hiccups before that point that we've got to deal with or that, I don't even know of saying we've got to deal with this right way to think about it, but that humans right now own these systems and they're available. The cost of accessing them, if you have the knowledge, is relatively low. So we think about the ability to use artificial intelligence for whatever it means you'd like, something akin to

the ability to hack, right? These are very similar kinds of skills that really are just part of now the smorgasbord of the tool set of governance, both official, right?

James Caton ([16:19](#)):

So governance at the national, state and local levels. We read about this with the policing. But also in terms, perhaps even in our response in terms of informal governance or we can ask the question, how is it that AI is impacting our communities? What does a community look like now that we're all on social media? And social media, a lot of our interactions are being driven by the social media algorithms. A lot of what was once private is now public, right? We never used to have, and I want to return to your question, but I think this ... I'd be doing a disservice if before talking about the big jump to singularity, we didn't think about the big problems we have in front of us right now. Right?

James Caton ([17:14](#)):

So, we have this problem, or not a problem, but a different feature now in our interactions where we have a very public life and we have very public social media life, we've lost a significant amount of privacy and we have all this data out there ready to be used to analyze us. Whether you're a corporation or really it's open to anybody who would like to use it if it's available. And there's plenty of speculating about who could be advantaged from being able to analyze your data, for example, or having a history of you that's very public. Right?

James Caton ([17:50](#)):

But if you wanted to, you could imagine like the job of being a private investigator, right? All that information about ... in your history, if it's not protected, could be used to compile a profile about you. We usually think about these sorts of things totally anonymously. But I'm skeptical that we should only think about them in terms of anonymous indicators about your purchasing profile at Target that was linked and dealing to your credit card or something like that. But rather, you have a Facebook account, you have a Gmail account, your information is all over the internet and you can build programs which are presumably artificial, and artificially intelligent, but you can build programs to gather all of that.

James Caton ([18:38](#)):

I mean, these are the sorts of problems we're facing right now when it comes to development of technology. AI is just another ... the more powerful AI is, I think the more powerful of a tool it is to be used in these governance processes, but also in these informal processes, right? And you can think of it in these terms in the same way that private businesses spy on one another, they sometimes are caught doing that or government spy on one another, right? All of those tools are available for people who have the skills and have significant enough motivation, oftentimes times monetary or power motivation, to collect that information. Okay.

James Caton ([19:18](#)):

So these are things that we tend not to think about because you or I or some normal, you could think of people living their everyday life, who don't think of themselves involved in the middle of power politics. Don't worry about, because if you're a good person, you don't create trouble, then what's it matter if people spy on you? But at the same time, in addition to the concerns about government spying and NSA and these sorts of details, that problem doesn't stop at the level of government. And I think this is something, as far as I can tell, it isn't very well appreciated. And I think that it informs our conversation about the potential dangers of advancement in AI.

Anthony Comegna ([20:03](#)):

Yeah. You're taking me back to a book that I edited and helped to republish when I worked at Cato by William Godwin, who's a founding father of individualist anarchism, libertarianism and socialism, all three. And he was Englishman, husband of Mary Wollstonecraft, Mary Shelley's father. And he wrote this book toward the end of his life called Lives of the Necromancers. And it's all about the history of a cultism essentially and different people in the priestly class and the governing class, their attempts to master knowledge and to use that to their own power and their own aggrandizement. And that over time, standard knowledge over the forces of nature translates into people's belief in some sort of mystical powers of the monarch or of the priest or priestess, right? And people believe in mysticism rather than science. But Godwin, coming out of the enlightenment, father of the woman who wrote Frankenstein, he's convinced that all of this is just simple scientific knowledge.

Anthony Comegna ([21:21](#)):

It's bits of flashes of insight into the way the world actually works that people monopolize and use to their own advantage. And entire civilizations get erected around protecting these particular individuals' monopoly on the power that science gives them. And it's full of these rich stories of the quest for essentially robots moving statues, machines that could think on their own, mechanical recreations of the brain, elixirs of life and all sorts of things like that. But he ends it being very optimistic in thinking that science is now open to everybody, that the day of belief in witches is over. And we have busted up these monopolies on technology and advancement and knowledge and that's a good thing.

Anthony Comegna ([22:14](#)):

But now, especially having talked with you about this a bit, I'm thinking are we returning to the days of witches and necromancers where the kings of the world, people like the president who has a nuclear arsenal at his disposal that could destroy the entire species, do they hold an immense amount of power monopolized to themselves? Should we be more concerned about the monopolizers or the fact that there are opportunities for demonopolizing some of these destructive technologies?

James Caton ([22:48](#)):

Yeah, I think every change or especially every significant change is both an opportunity for good developments and bad developments. Some changes carry with them opportunities for more bad developments or more chaos than others. When we talk about these technological advancements and changes in our understanding, and social positions that arise with these things, we delve right into social contract theory. And what we have to think about with that is how quickly can we withstand social change that transforms our positions within the hierarchy that we exist in within society.

James Caton ([23:36](#)):

So, James Buchanan thought about this pretty extensively, but we're continually bargaining over our positions in the hierarchy in terms of the social contract or continually bargaining over what our powers are within our positions. What's okay for us to do? What's not okay for us to do? And one way that we bargain is through politics. We just had the Iowa caucus that was one means of speaking of technology and problems that may occur with technology. That's one means of bargaining over your position. And what this technology may be doing. It could be empowering particular persons. And I think this is what you're saying, it could hyper-empower persons and organizations with the special knowledge that they are collecting and developing.

James Caton ([24:27](#)):

And the question is, is how does that shake off in terms of the structure of society abroad, in terms of international politics, in terms of the national politics, how is this going to change things? How does this change things at the local level of the community? Is it going to impact the sorts of modes and rules that we've been following in our everyday life with people that we see everyday or that we've known for our entire lives? If you live in a community like that. We see transformations, and this is part of what I think what Newmann's question was when does the change start happening so fast that the social structures that we're now using to navigate and coordinate with each other are no longer keeping up ... we're not able to maintain the roles that we had before. Does this lead in our conversation to hyper-empowerment of particular agencies? Is it that we're going to live in a world of Amazon feudalism?

James Caton ([25:37](#)):

I'm skeptical when I say that, right? I'm not saying that is the world that we're going to live in, but I think these are the sorts of questions that we have to be willing to grapple with when we think about the effects of technological change, because it's permeating every inch of society, it's radically changing our lives and the ways that we communicate. I mean, even if you think about, at the level of grade school, it's radically changing bullying and empowering, not just bullying, but empowering collective bullying where nobody has to take responsibility for this sort of thing.

James Caton ([26:09](#)):

This is happening and we have to find modes of learning to protect ourselves, each individual. We have to find modes of finding communities within that, where we don't feel like that's a problem. If I have to find personal empowerment in the process. None of these problems I think is an absolute problem or else you'd have to be a pessimist to say that the world is not going to be good at all. But we have a lot of new problems that we have to confront because of technology and it'd be difficult to make a finite list of the number of problems that we're going to have to confront because of it.

Anthony Comegna ([26:51](#)):

James Caton is a Professor and Economist at North Dakota State University, and he holds a PhD from right here at George Mason. He has a long history with IHS by now, and we were just thrilled that he could lead such an important and concerning discussion for us. If we whet your appetite for the bizarre, troubling and possibly glorious future, tune in next week for more ideas in progress.