

Anthony Comegna ([00:19](#)):

Last week, professor James Caton and I sketched out the broad picture of what technological change might look like over the next 30 years or so, but what about that balance of doom and gloom with liberal optimism? Is there room in such a tumultuous time for liberal futurism? What would it look like, and what can liberal scholars do to ensure that the future is freer than the world today?

Anthony Comegna ([00:46](#)):

Okay. Now the reader for this program on artificial intelligence and liberal futurism, it is admittedly a bit heavy on the doom and gloom side of things, and that's my fault, I suppose. I pick the readings, and I have a morbid fascination with the doom and gloom arguments and prognostications about how awful the future might be. But that's mainly because I want us to be able to reckon with that and manage in the intervening years to construct a real liberal futurism out of it. So I'm wondering first, to start this episode, can you give us an overview of some of the articles or authors or arguments here that really leapt out to you as a specially pessimistic?

James Caton ([01:34](#)):

The future is bad enough, but there were some, we're actually looking at the problems of the present. And the present is frightening. It shouldn't have been a surprise to me, but again, there's so many kinds of problems in the world that it's impossible to think about all of them, to be cognizant of all of them, and then well better for us to be depressed about all of them. But the piece by Garvie, Alvaro, and Frankle on the facial recognition by local police and state police, that really took me aback. The idea here is that we're being watched all the time. So as you said earlier, not only is there plenty of public information about you out there for people to navigate if they so wish, including the police, but your face is public information and your existence in public is public information, and there are cameras everywhere.

James Caton ([02:48](#)):

This made me a bit uncomfortable to think that we have facial recognition technology and we have cameras that are explicitly intended for that. Are you think about, "Well, what about the number of cameras that probably aren't intended for that," but at some point in the future, especially a less liberal future could be used for that. I think that is a really frightening idea that every time you step outside, you could be recognized and then you could be considered maybe a suspect in a crime, because your face looked similar to somebody who committed a crime.

Anthony Comegna ([03:28](#)):

The racial component of that article was especially troubling to me.

James Caton ([03:32](#)):

Yes. This is really a problem with AI is it can be very good at identifying some things, but there's always going to be an error. And speaking of rationality, every rational system has errors. Human rational system have errors and artificially intelligent rational systems have errors, and depending on the environment, they can be better or worse. It so happens that the ability of these programs to successfully detect the features of darker faces is it's more challenging for these programs than for lighter faces. So you get more false positives, and that is a problem, a very real problem. It might not be if we had a little bit of humility about the reaches or abilities of artificial intelligence, but I think just

being human, the more natural inclination for those not predisposed to be suspicious of their own interpretations or of the tools that they're given is to say that the AI is right or it's right often enough that we should use it for this facial recognition.

James Caton ([04:46](#)):

But if I was somebody who got brought into the police department because my face was falsely recognized as a positive, I don't care if it's an hour or a day, that is psychologically traumatic, because now you know that any point in time you could end up being arrested because of this error. And this is something that makes you feel unsafe in a liberal society where presumably we want to feel safe if we are not not criminals. We don't need this extra element be pulling us in, and for which right now it's highly ambiguous and maybe even a little bit... I'm pessimistic about what kinds of constitutional protections we actually have for facial recognition in public.

Anthony Comegna ([05:34](#)):

Yeah. It's very, very troubling, and it brings me back to a point you made in the last episode about the real problems are the even more fundamental issues first that we still have to work on and work out before we get to this new stuff that's constantly piling up. So we have to obviously have more intersectional understandings of the systems here and how they're constructed, how they're implemented, how risks are assessed given what the facial recognition system has produced its human officials. So everybody along the chain has to have some sort of racial sensitivity and awareness of how the intersections of a person's life and identity compound together and produce totally different situations for each person. And I don't know that there's much understanding of that or ability to practically incorporate that into things like the police surveillance state.

James Caton ([06:44](#)):

Yeah. And I think, especially simpler knowledge, does a fair job of abstracting away from details that may be in a certain situation aren't pertinent. But this goes on to expert knowledge, as well, and our trust in experts. I was really bothered in thinking about these problems with artificial intelligence, that the natural inclination for people to trust experts could leave us in a really problematic situation in terms of not only this facial recognition, but let's say, for example, a artificial intelligence selecting the sentencing for a crime, which is in conversation right now. This is not 10 years from now. This is a topic of conversation today, and just like we need to have suspicion of experts and expert interpretation, we need to have suspicion of this AI for whatever purpose we want to use, especially when it's life and death situations like this.

James Caton ([07:50](#)):

The problem that we're facing in terms of AI is that most people aren't experts in this stuff. It's very difficult to be an expert in AI, and if you're an expert only in AI, even then we probably shouldn't trust you. It's like Hayek's comment about an economist that only works with economics or only considered himself an economist is particularly dangerous. And so we have these feuding impulses or impulses that feud with the reality that we're facing. On the one hand, people want to trust experts, but on the other hand, to be an expert in AI and its applications, it actually requires being not just an expert in AI, but at least being sufficiently knowledgeable in many other areas, especially in terms of, say, political economy. What are the effects in terms of political economy, the distribution of power, the effects, therefore, or as part of this on our political systems, our political organization, on individual liberty is a part of this.

James Caton ([08:58](#)):

But my interpretation is that there's so much excitement about these uses and so much opaqueness about their application or in their application, who's using it for what, and how is that affecting you that those who seem to be sprinting ahead, and this is reflected by the article on facial recognition. They're sprinting ahead without much encumbrance.

Anthony Comegna ([09:28](#)):

Now, I think that perhaps the most pessimistic article out of the whole bunch here though, it has to be Nick Bostrom's article on black ball technologies. And a lot of people have criticized Bostrom and his clique or wing of this scholarship as being very overly doom and gloom and worrying about scenarios that are very unlikely to ever come to pass. And his argument is, "Well, if they do it all, then everything's lost. So we should be concerned about it and we should be sure, especially our ethics line up with the kind of future that we want."

Anthony Comegna ([10:12](#)):

And now, he has this article on a so-called black ball technology, which refers back to ancient Greece. If you pick the black ball out of the bucket or whatever, then, I don't know, you're the one who has to go do the most uncomfortable job or something like that. You get banished from the town or whatever. They had different traditions drawing black balls. And his idea is what if we have a technology that turns out to be so easy to produce that it just spreads everywhere, and it's so destructive that there's no stopping it?

Anthony Comegna ([10:46](#)):

So for example, he says, "We got very lucky that it's actually incredibly expensive and difficult to produce nuclear weapons," because what if it was really easy? What if all you needed were two sheets of metal, a pane of glass and a wire to connect the sheets of metal together, and boom. You get a nuclear explosion as soon as it's all hooked up. If it was as easy to create a nuclear bomb as it was to create a battery, the Romans would have destroyed the world, because they had simple batteries. And what if we're on... The whole thing is Kurzweil says, "Well, the singularity is going to be marked by computers that you could buy for \$1,000 that will have the intelligence of the entire human species." Is that a black ball technology? What the hell happens to the world after a change like that? I certainly don't know. I don't think anybody really knows, but what are the kinds of dangers that we face given the knowledge we have now about what's coming up?

James Caton ([11:47](#)):

Well, I think a lot of these questions come back to governance, and we have some major problems if we continue this trend of increased empowerment of the central government relative to local governments, increased empowerment of government actors relative to local actors, because the real problem is that in the liberal society, you technically don't have to be part of a community, which means you don't have to be monitored. If you're in a community, you have more monitoring within the community than if you're not, although communities also monitor outsiders, this sort of thing. But it's possible on the liberal order, you could go live in New York, get an apartment, work 9:00 to 5:00 and not really see anybody there. It's totally possible, not bother anybody except go to work, and you could even work remotely.

James Caton ([12:49](#)):

And so you don't have to be part of community. You can support yourself in the modern society without having what Buchanan would refer to as a being part of a moral community. It's not just knowing people and seeing the same people every day, but actually participating in the same sort of metaphysical entity called the community. You all share it in your minds, and you all share your positions and how you interact in your minds in the community. And so within a community, there can be governance of these kinds of things, because most people know about what's going on with anybody in the community. I'm not say that there can't be secrets or can't be asymmetries of information, but there's a lot of incentive participation. Or community membership has a lot of value, and you tend not to do things. If you value the community membership. That would be highly damaging. The danger is what if we find somebody who doesn't value their membership in society? They're not part of a community. And so one day after getting fed up with their work, we see this all the time with mass atrocities. Well, what if instead of a gun, it's a nuclear weapon? I think this is the question we're asking.

James Caton ([14:13](#)):

We have a higher likelihood of that the less of a job local communities are doing in governing interactions amongst themselves. I think the less incentive you have, less disincentive you have for something like mass destruction. You've got to value where you're at. And if you don't have any connections in the society, then you don't necessarily value where you're at. So we have a lot of community building to do if we really want to take this problem seriously, because I'm really skeptical about other solutions, because they themselves... He talks about empowerment of central government to be able to maybe limit technology or the use of technology.

Anthony Comegna ([15:02](#)):

International government, too.

James Caton ([15:04](#)):

International government. It makes for a terrible sci-fi movie. We might as well let's just implement Skynet now, and then we'll just get it over with. But they themselves, and he recognizes this, that solutions like that tend to themselves be something of a great filter. They could prevent us from actually progressing and lead to a reversal of the prosperity and freedom that we've seen. So those aren't solutions. But unfortunately, I don't see Bostrom proposing local solutions. I don't see him thinking about... And this is where I think it's really important to be able to have some idea about social science and how communities work and how individuals cooperate together. We need to know enough Elinor Ostrom, so to speak, to think about how we can use principles of local governance to integrate people into society so they don't have the incentive to do that. Now I, of course, have not dealt with the problem of an accident, which is the next issue. But at least as far as mal-intent goes to prevent issues in that, I think that will be very important moving forward is increased participation in membership in your communities in a manner that enables internal governance.

Anthony Comegna ([16:29](#)):

Yeah. It struck me as very strange. If what you're concerned about is one of these black ball technologies that is just unexpectedly capable of producing mass doom, and you're worried about the fact that it might just pop up out of nowhere. It could be in the proverbial terrorists cave, or it could be in somebody's basement, like you said, somebody's disgruntled with their job and their go nowhere life or whatever. And it could be any number of government laboratories scattered all over the world from our hundred plus different national governments that are around. It could be all sorts of different things. So

if you're concerned that the problem's very nature that it could pop up anywhere very cheaply, very easily, why in God's name would you not recognize the intense complexity of that problem and try to solve it from the bottom up as opposed to the iron grasping hand from the top down trying to control everybody and make sure they never produce a black ball. That does not at all seem possible nor intuitive to me.

James Caton ([17:45](#)):

As an economist when you think about these sorts of problems, they're inherently ethical problems and they're ethical problems that a central government is not prepared or basically should not be trying to deal with. The solutions definitely need to occur locally, but we start moving from the realm of economics to the realm of ethics here in that action becomes much more likely to be harmful. And I think that a lot of traditions deal with this sort of problem, understanding that not all action that is, obviously, not all action is good. This is the foundation of law. This is the foundation of a lot of religious organization.

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

We need to put limits on all persons. Even the King we need to put limits on. This is a legal issue. This is an ethical issue, but I don't see any way around this except ethics in terms of people recognizing that every action that they take, every development of a technology that we get can potentially be a black ball. Every use of a technology can potentially be a black ball pulled from the urn that not only gives us a terrible job, but and everything for humanity, or at least ends the free society that is responsible for the progress that we've had in particularly over the last two or 300 years. So this means living in a world where trying to be ethical may just be acting less, may be not pursuing value, pursuing goals, but being extremely self critical about what the second order effects of your action can be, and thinking that far out.

James Caton ([19:43](#)):

And probably everybody can't think in those terms, but we definitely need leaders in communities who are thinking in these terms about, "Okay, well if I develop this general AI, what's at least a couple of degrees out that I can expect all these different scenarios?" At least I can build a game tree in my head and say, "Okay, well maybe I shouldn't develop it in this way, or maybe it shouldn't be made accessible to everybody, and we need a way to keep it secret." We've got to be careful.

Anthony Comegna ([20:19](#)):

Is there no secret libertarian bullet, a silver bullet here that will just magically solve all of these problems for us? Will spontaneous order somehow save us from a hellish future of dystopia? I used to argue with some of my friends, for example, that "Well, AI's would be so intelligent that they would, of course, recognize the benefits of comparative advantage, and we always have a value then even if they could do literally everything better than we can, there are still comparative advantages to be found, and there's space to live together. Surely they would know that. They'd understand that they'd leave us alone. We can help them some in some small way."

James Caton ([21:06](#)):

Well hopefully, a general AI, if that was ever a thing, would recognize that. But we could ask the same question about aliens. If aliens approached us, would they recognize comparative advantages is the goal, or would they just see our planet as a good planet to colonize without much care about the value

of human life or the quality of life for us. When you say libertarian principles, so I think competition, for example. A lot of times, we just fall back on competition. The market will work it out. That's not a particularly libertarian principle, but I think it is akin. It is akin to libertarianism. It's one of the fundamental spontaneous orders that we see that are good. But I don't know why, and this is why I like to hear your thoughts in asking the question. Would we expect that the spontaneous orders developed around the new AI technology would generally be good spontaneous orders? Is it possible to get bad spontaneous orders? And if so, what can we do to deal with that?

Anthony Comegna ([22:20](#)):

Well yeah, that is actually why I asked, because I think it's a conundrum. It's a quandary and we have to recognize that, yeah, a lot of times people want godawful stuff, and they get it because markets do actually work pretty well at delivering people what they want even if it's terrible. If ethically it's the wrong thing to do or to want, like slave trading markets and the whole political economy that in society civilization that developed around not even maximizing productivity or maximizing wages, but maximizing white supremacy, really. And I think it's a problem that markets work sometimes, and it's a problem that we get what we want through spontaneous orders, because sometimes we want some godawful stuff. So I, like it sounds like you do, I see so many of these issues is fundamentally about individual ethics.

Anthony Comegna ([23:24](#)):

So to me, I think people sometimes do lean too heavily on what... It's a caricature of classical liberals that we just think the market will work things out, that it's some disembodied force out there in the universe. But sometimes I think we do slide into that, at least we start hoping that world works that way and that things will just work out well over time, that the hockey stick is unending and we'll always get progress, but progress doesn't always look good. Change might be inevitable, and the different orders that we get might be spontaneous and driven by consumer demand, but it's sometimes awfully messy.

James Caton ([24:10](#)):

In this respect, this is where I felt a lot of skepticism. I feel a lot of skepticism in looking at the pronouncement of Kurzweil. These ideas are a significant part of the reading, even if this isn't a reading particularly about his ideas in any way, but the name comes up quite a bit, and we talk about... So Kurzweil makes a prediction, a couple of predictions. One is the Turing test actually being passed by 2029.

James Caton ([24:46](#)):

So those who don't know, the Turing test is a test where a human talks to a robot. You could say a human judge talks to a robot for several hours. This is at least the circumstances that his bet is made on, and this human being will not be able to tell the difference between the other human that they're talking to and the robot that they're talking to. And if that's the case, then we can say that the Turing test has been passed, or we have moved to this new stage of AI at least in terms of natural language processing we certainly would have with that.

James Caton ([25:29](#)):

If you listen to him, he doesn't just make predictions about the Turing test. He doesn't just make predictions about the singularity. We hear predictions about our biology and predictions about tech applied to our immune systems, and abilities of uploading consciousness to, you say, the Cloud by

nanotechnology being able to basically infiltrate your entire nervous system. These are concepts that he's talking about. Exploring, you're always going to have fascinating ideas versus things that actually seem realistic. But I don't know that we should always be optimistic, if we can imagine these worlds, that we should always be optimistic about what that's going to look like. And I'm just very skeptical about the pace of development of technology he's expecting alongside relatively rosy outcomes, because he is an optimist, and he's openly an optimist.

Anthony Comegna ([26:38](#)):

What then to your mind for all the especially graduate students out there either about to attend our colloquium tonight or who are trying to figure out their own research agendas here, how would you translate this phrase liberal futurism into a practical research agenda?

James Caton ([27:00](#)):

I think this has to be a local liberalism. Again, Buchanan makes this point that the liberal society depends on strong local communities and coherence between those communities in terms of liberal expectations. And we have lost a lot of the element of community in the United States. We are losing a sense of what it means to be American I think partly as a result of losing a sense of local community. He gives a comparison to say Japan is having a medic community, because they have the common culture and common expectations. If you have local governance and local development, local oversight essentially from all these different scales, every network you're involved in, every in a local town. But a community isn't just a town. It's every organization that you're a part of and how that links to every other organization.

James Caton ([27:56](#)):

If we have strong local communities with strong, well develop moral norms, we might maybe be a little more careful about these developments, or we might ask more questions about the second order effects of the sorts of practices that we're engaged in. So if you are a graduate student, I think research on local governance is a hot topic. Research on governance of AI maybe within corporations, maybe between corporations, agreements that can be made between different groups who are working on all of these developments, that can be forward looking, that don't have to be coming from the state, don't have to be X anti-regulation that is strangling all of the great potential we could get from this, but rather is akin to something like an underwriter's laboratory except for AI.

Anthony Comegna ([28:55](#)):

Now here's a surprise bonus question for you. The scientist is in the research lab. He's got the first legit general AI. He's about to flip the switch and turn it on, and you burst through the door and you say, "Wait, wait, wait. Hold on. Before you turn it on, have you," and then what do you say? What's your question for him? "Make sure you've done this before you flip the AI on."

James Caton ([29:23](#)):

Have you made sure that it's contained within a system not connected to the rest of the world?

Anthony Comegna ([29:30](#)):

You're really worried about the Skynet problem.

James Caton ([29:32](#)):

Let's start there. You can have general AI in a bubble, and then maybe we can talk about different linkages or developments from there. But I'm actually really skeptical. If you had a general AI, and it was let out into the world, so to speak, and it's existing on all these different systems, every program has bugs. A general AI would be another program. So I'm skeptical that you'd instantly have I-robot. I don't think that that's the immediate outcome with a general AI. It's more like you just have a really smart robot dealing with the world.

Anthony Comegna ([30:29](#)):

Of course, you'd say that. You're a Cylon. So.

James Caton ([30:32](#)):

Well, I didn't want to give that away.

Anthony Comegna ([30:38](#)):

Well, there you have it. Now, how's that for a research agenda in liberal futurism? Sure, the means of coercion are expanding rapidly, but we can do everything possible to maximize the means of liberation at the same time. But we've been here before. We've had other technological singularities with mixed records of progress and pain. And the future's course always depends on our actions here and now. The future may never be the result of human designs, but it will always be the result of human actions.