SCIENCE OF THE NOOSPHERE

Athena Aktipis and Michael Levin

with

David Sloan Wilson

David Sloan Wilson: Welcome Athena Aktipis and Michael Levin to this episode of the Science of the Noosphere series. This series is inspired by Pierre Teilhard de Chardin and his concept of the noosphere, which is like the thinking, the mental dimension of humanity, which he said originated as tiny grains of thought with small scale society and then was enlarging to ultimately encompass the entire earth as some kind of global consciousness. The purpose of the series is to actually place this concept on a solid scientific foundation and that maps on to a concept that you know well, major evolutionary transitions in individuality. So you might or might not know much about Teilhard, but you know a lot about major transitions. We'll be focusing on the single cell and multicellularity as those particular transitions. Well, let's begin first of all having you introduce yourself as people, how you wandered into this kind of work, and how you encountered the work of Teilhard, if at all. Athena, why don't you begin and then we'll pass to Michael.

Athena Aktipis: Well, for me, I've been just fascinated by the idea of major transitions, really from the time that I was in my late teens. This notion that you have levels of organization where you get regulation of cooperation, and that allows you to build additional levels, and additional levels, has been really a foundation of my thinking from the very beginning. The notion that there is this layer of collective cognition that's happening in humanity, to me, it's a reflection of the fact that under all of these different levels you have computation and information processing happening in surprising places. Sometimes it's decentralized, sometimes it's more centralized, but you do have this sort of unexpected prevalence of information processing happening everywhere.

DSW: Yeah, totally. ML, how about you?

Michael Levin: Yeah, well, I guess all of my work probably stems from two kind of foundational experiences that I had as a kid. One of those was looking inside the back of a TV set. This is like one of those giant old school TV sets with the vacuum tubes. Thinking about the process that led to somebody knowing how to put all those things in just the right configuration to have something interesting happen. And conversely, playing with bugs and insects outdoors, asking basically the same question, right. Now you've got a different type of creature that probably has some internal perspective, and goals, and preferences and so on.

Just asking these really fundamental questions about what it means to have some sort of a body that is made up of parts, but nevertheless has a kind of a centralized, Is the subject of collective memories, and preferences and goals and so on. So my background is originally computer science. I got a an undergrad degree in computer science trying to work in artificial intelligence and then did biology hoping to really understand the decisions that cellular collectives make as a morphogenetic process, viewing morphogenesis as the behavior of collective intelligence.

AA: Michael, while you were staring at the back of the TV screen, I was staring at the front of the TV screen. When I was four, I would absolutely...

DSW: Like most people.

AA: Yeah. Well, there were two things I was obsessed with. The first one was that movie, Annie, because it was just musical amazing choreography. The other was the NOVA special about morphogenesis. I had watched that over and over and over again. I think that's really where my fascination with emergence and multicellularity really started was just watching that over and over. I was asking my parents how does it do that? They gave me vague answers, all like chemical gradients and gene expression. But I was like, how does it know to do that? Definitely spent many, many hours when I was four, watching that over and over and over.

DSW: One memorable thing that Teilhard said was that, although in some respects, we are just another ape species, little different than other apes. In other respects we're a new evolutionary process — cultural evolution. That makes the origin of our species as significant in its own way as the origin of life. To take that seriously requires, actually, knowing a lot about the origin of life and biological levels of organization in order to compare human levels of organization. When we get to the single cell, many people have thought that that's relatively simple, that the history of life has been one of increasing complexification.

And I think where your work stands out is in demonstrating that no, the complexity of a single cell and the first multicellular organisms is absolutely amazing. Life itself in its first forms has an extraordinary form of complexity, including a form of mentality. A form of mentality, and even a form of group level mentality. So a lot that's associated with the noosphere in humans can be understood at the very initial stages of life. I've stated this very crudely, you guys are the experts. Michael, why don't you begin just elaborating on that theme, the complexity of the single cell and then, Athena, take over. So Michael.

ML: Maybe I'll just sketch a little bit how I tend to think about these things. One thing that I'm interested in is trying to come up with a framework that is able to handle in a very practical way, in terms of empirically driving research and so on, truly diverse intelligences. I'm interested in what is the invariant between all intelligent things, no matter what their embodiment is, no matter what their providence has evolved, designed, exobiological, artificial, whatever. In doing that, this is a framework I've called TAME T-A-M-E for Technological Approach to Mind Everywhere. What I take to be central about intelligence is problem solving, in some particular space. In particular, this idea that we recognize agents or selves or individuals and we can compare them, no matter what they're made of or how they got here, by the scale of goals that they can pursue.

To me, the fundamental building block of all of this is some sort of goal-directedness. It can be as humble as a homeostatic loop keeping a metabolic state, which even bacteria can do or it can be very large scale things. The spatial-temporal scale of the goals that you are capable of pursuing, right? It may be very small in case of a microbe, local concentration of sugar or something, or it could be very large in the case of humans. We could talk about the transition to the very special transition to humans in terms of the goals that they can pursue is what to me is central. When you look at it that way, in fact, we see evolution as a process that can scale cognitive selves from very humble local goals to large scale goals first, like making an organ, making a whole body. And then, of course, behavioral goals in three dimensional space.

We see a very smooth process of scaling these goals throughout evolution. I see bodies as these overlapping nested communities of selves with local goals in various spaces, in physiological space, transcriptional space and metabolic space, of course, anatomical morphospace. To me, morphogenesis, like developmental morphogenesis regeneration, those kinds of things, are the process of the collective intelligence of cells operating in anatomical morphospace. I think intelligence defined operationally, that way goes all the way back to the beginning. What we see it's scaling through evolution.

AA: Michael, I'm curious, just listening to you talk about this idea of nested goals. I can't help, but think about human groups and the challenges of getting lots of people on the same page with goals, or even if people have the same goals, getting them all pointed in the right direction at the same time, almost like thinking of an organ trying to develop, right. It's not enough for them to all have the same DNA. Everyone has to be in the same expression or in the right expression state, right. In order to accomplish a shared goal. When you think about humans having goals together, do these same principles apply or are there things that change when you go to the level of thinking about humans.

ML: Yeah. Very interesting question. I think looking for scale-free principles that might carry over from subcellular networks, through cells, to human societies, collectives, ants and bees and these kinds of things, is very important. I'm in no way an expert on societies and human activities, but I think we can see something very interesting happen in the cellular case. I think the cellular case is very instructive for how this scaling happens. I'm really looking forward to talking to you about the cancer aspect of this, because I think it's actually very foundational there. What will happen and this is what we've observed is that you have individual cells, which are themselves competent in their local spaces. So they solve transcriptional problems, physiological problems, basically homeostatic units. They're trying to keep certain things within parameters. Evolution discovered a really interesting thing fairly early on, which we call a gap junction. This is an electrical synapse. The thing that's magical about these gap junctions is that if you imagine traditional signaling, extracellular secretion, I secrete some sort of chemical and it goes over and hits this other cell. When you have two cells, one signaling to the other, because these signals come from the outside, the recipient cell knows very clearly that these originated from the outside, so it can ignore them. It can do something about them. It can remember them, whatever. When two cells couple to each other with these gap junctions, what happens is that these signals propagate to some extent directly into the intracellular milieu from one cell to another.

What happens then is because these signals don't have any kind of metadata on them that says, where did they come from, the ownership information is wiped. So when you have two cells that are tightly coupled by gap junctions, one cell let's say gets injured or has some new experience, it generates a calcium spike that appears in the other cell. That second cell cannot tell whether this is its own memory or some other memory of some other creature. Basically, this connection by wiping ownership information, that kind of connection creates an almost mind meld where you can't tell where your memories and some other memories begin. When this happens in a large network, you end up making these electrical networks that have the computational capacity to measure, remember, and store goal states that are huge that a single cell couldn't possibly...

DSW: We're talking bacteria here just to let everyone know. Some of what I've read of yours, Michael, talks about biofilms basically. Bacterial biofilms as superorganisms. Among the superorganisms, we all know about the eusocial insects, but now we have these microbial biofilms, which really do qualify as superorganisms complete with the analog of nervous systems. Talk to us about that just to nail down this idea of the complexity—before we even get to nucleated cells—of bacterial biofilms. Tell us some stories about that.

ML: There's some beautiful work from UCSD looking at electrical signaling in bacterial biofilms as very brain-like. This isn't our work. I think because evolution discovered very early on that electrical networks are a great way to scale and propagate and coordinate information to store memories, they're just very convenient. It's no accident that the nervous system uses them, that we use them for computer technology. In the body, what happens is that these networks are now capable of storing much larger types of goal state information. The individuality of the components is wiped and erased to an important extent, because now you have memories, you have goals and activities that are property of the collective that individuals couldn't do.

So just to get back to Athena's comment about humanity, I think from the perspective of coordinating morphogenesis and avoiding cancer, these kinds of gap junctional connections are really critical. One might be tempted to say, look the way to scale goals and make sure that everybody's on the same goal is to mind meld this all into this larger scale network where all those can take place.

I think we have to be very careful about that be because when was the last time you were too concerned about what happened to your individual skin cells, right? We shed them all the time. We don't really worry about it. You might go rock climbing or something and tear a bunch of skin off your hands and you feel good because, hey, I went rock climbing. The individual cells may not be happy about that at all. Presumably there's some sort of optimal linkage mechanisms to be found that allow us to reap the benefits of a larger collective IQ while at the same time holding on to the respect for the goals and the needs of the individual pieces that go in there.

DSW: Well, we'll return to that later on. Athena, back to you.

AA: Yeah. Just hearing you talk about the gap junctions, it sounds just like you've got someone under your skin or like you care. You have a lot of empathy and you can't help but feel the feelings when you see someone else feeling them. Right. It's fascinating to then think about we do have different problems that we have to solve as human groups and different constraints, right. We're not all genetically homogeneous. There's possibilities for cheating. They have to be regulated in a much tighter way than you have in a multicellular body. Anyway, David, I'll stop because I know you want to bring the conversation back around to this stuff later, but I think it's a really rich area to explore some of the similarities at these different scales.

DSW: Totally. We don't have to wait. Large scale society has a lot of the same characteristics, especially the concept of the invisible hand—that individuals can be operating locally, maximizing their local utilities, and somehow it works out for the benefit of the common good and the individual does not have to have the common good in mind. What would that be, but like a liver cell operating according to its local environment, but nevertheless, the whole thing coordinated so that the common good is benefiting.

AA: I think the question is are the rules that the individual is operating by, whether they're a human or a liver cell. Have those been designed and calibrated for higher level function or not. Right. A liver cell... Yeah, they have been designed and calibrated for that higher level function, but in the case of humans, there's conflicting goals. There are some that are for higher level function and then there's others that are for more individual level and negotiating that I think is one of the big challenges of applying some of these principles from multicellularity and regulation of multicellular bodies to large scale human systems.

DSW: Yeah, totally. Just to... Go ahead, Michael.

ML: Well, I was just going to say, I think this is a really interesting question about being calibrated to specific goals, because one of the things that we've been working on, and of course, other people had as well in the past is the diversity of behaviors that can be executed by cells with exactly the same genome, and the plasticity with which cells of any given genome can solve problems in novel environments that they haven't seen before. One of my favorite recent examples is this thing we've been doing we call xenobots. These are novel proto-organisms that never existed on earth before. Basically, what we do is we take skin cells off of a frog embryo and we don't add anything. We don't edit the genome. We don't add any novel transgenes or nanomaterials or anything like that.

We take away something. We take away the other cells that are normally instructing and constraining these cells to have this boring two dimensional life on the surface of the animal. We put them in a new environment and we let them reboot their multicellularity. We ask them in effect, what, in the absence of these other cues, what do you want to do? What is your default behavior? We find out that actually they self-organize as many things they could have done, but in fact, they self-organize into this novel shape with novel behaviors, novel capabilities with self-directed behavior, in fact, that raises some really interesting questions about what did evolution actually do when it developed the frog genome, because being an excellent frog is one thing it did, but that's not the only thing it did. In some way the exact same hardware in a novel environment is able to now follow different kinds of goals, including figure out a novel way to reproduce itself. They reproduce by kinematic replication instead of the normal froggy reproduction way, which we've made impossible.

So I think there are some really profound gaps in our understanding of what are things actually optimized for, right? Because there's a tremendous amount of freedom and creativity when you release some of these constraints. You find out that there's actually this agential material that evolution has to work with, these cells. If you do nothing, they're going to do things on their own. The best you're going to do is guide them to various adaptive outcomes. I think that's a very interesting set of questions.

AA: Michael, do you ever feel like a mad scientist sitting there? You're like, I wonder what these cells will do if I take them out of the frog and put them in new environment?

ML: Constantly. Yeah, that's my preferred state of mind. Yeah, absolutely.

DSW: Well, didn't you get an eye to develop within the gut of a tadpole or something like, I think?

ML: That's exactly right. We should talk about that. Yeah. We've been playing with understanding the cognitive medium that holds the patterns to which these collective intelligence is built and not terribly surprising, it's bioelectrical. It's like a non-neural version of neuroscience basically. In certain cases we've learned to read and write that electrical information, so we can overwrite the native pattern memories that these things have. Then you can go to a different group of cells in the frog that are otherwise going to make gut, and you can say to them, you should really be making an eye. In fact, that's what they do. They'll make an eye. In fact, what we found is that when you place these eyes, even on the tails of tadpoles, those animals can see perfectly well.

The eye cells will form even though they're sitting in the middle of muscle, instead of in their normal place near the brain. They form a nice eye. They make an optic nerve. The optic nerve runs around and looks what to connect with. Doesn't find the brain. Sometimes finds the spinal cord, so it'll connect to the spinal cord. Those animals, when they don't have any primary eyes, just one eye sitting on their tail, we can train them for visual cues. We've built this machine that trains them to stay away from a certain kind of light. They can see perfectly well.

AA: Are they evolving for evolvability in a way that other organisms aren't? You think there's big variation may be in this, like some organisms just have the ability to change up their body plans when things change.

ML: Obviously we haven't done this in every organism, but I have some examples in other models where I don't think this is special. I think this is a fundamental feature of life and I'll tell you why. The amazing thing about that kind of…and I call this whole business multi—scale competency, because what it means that evolution can do as a process is rely on the various sub components to get their job done, even when circumstances change. What it means is that if we're using an architecture where nothing is really hardwired. Well, except for metabolic things, things aren't really hardwired. What you're doing is you've got modules that are going to try to get their goals done, meaning homeostatically, they're going to try to achieve their objectives, their agendas as Dan Dennett calls it, despite the fact that their circumstances change.

This really potentiates evolution. It smooths the fitness landscape, because it means that if you've got some mutation that moves the eye to the wrong location, but it does something good somewhere else in the embryo. You don't have to wait until you find that that second mutation without disturbing the eye. It's fine. The eye will do its job. In the tadpole we show that craniofacial organs, if you start them off in the wrong position, we call them Picasso tadpoles. They will still move around in these novel paths and give you a perfectly good frog during metamorphosis. It doesn't matter that they're out of kilter early on.

So this means that evolution can really explore all kinds of things without immediately incurring the penalty of just the wrecking the whole thing. The minute you move something, just everything is wrecked. When you look at this in all kinds of species, right? Human embryos, you can cut them in half and you get two perfectly normal monozygotic twins, because it's incredibly plastic. Half the embryo realizes that the other half is missing, regenerates what's needed, and you're good to go. There's tons of these beautiful examples of plasticity and problem solving of this kind of competency. I think it's everywhere. I think it's what makes evolution workable basically.

AA: Mm-hmm. One of the thread then that goes between your and my work is this idea of beautiful monsters. I've thought about it in terms of the crested cacti. They have this cancer-like form, they're beautiful in the patterning that results. Your version is maybe a little more mad sciencey, but...

ML: Super interesting point, because understanding what a monster is and understanding what a birth defect is, right? Chemistry doesn't make mistakes, but biology, if you're going to say it's a mistake or an error or developmental defect in some way, it's with reference to an expectation that was supposed to happen. One of the interesting things that we see now is that these attractors in morphospace, different ways of forming bodies. We can take a planarian flatworm and you cut off the head, they of course, regrow the head. Normally 100% of the time they regrow ahead of the correct species. What we can do is perturb the electrical conversations between cells that help them remember what head shape they're supposed to make, and as a result, they will stochastically make heads that belong to other species of planaria.

There you have a complete mismatch. The genetics is unaltered. You've still got the same wild type genome, but the cells will land in a different attractor in morphospace. They'll make a head with a brain shape with a stem cell distribution, looking like planaria that are 150 million years distant, right, from the exact same genome. So this idea of what a monster is...and we can make tadpoles that have zebra fish-like tails and faces looking like other frogs and so on. This is speciation versus developmental defect. I think you're exactly right. It's very interesting.

DSW: So basically, the raw material that evolution is working with is much more agential than most people think. These lower-level units are organized into modules within other modules, and they're coordinated, and that results in this tremendous diversity that, if you know the mechanisms, you can create these monsters. So there's that, that's kind of like the variation part of a variation, selection, replication process. All along, we've been talking about the importance of levels of selection. That this amazing diversity that arises has to pass the survival test, basically. The survival and reproduction test. And so the selection at the level of the individual organism is now winnowing these. That's what distinguishes a monster from anything else, is whether it can survive or reproduce.

And so higher level selection is then operating on this amazing diversity that arises in this coordinated fashion because of these agents, basically in these homeostatic units and all of that. And right away, we could map that back onto humans. And the idea that individuals are like these lower-level units, they're tremendously generative, they're doing all of these things, but that's going to result in the social manifestation, which might work well or poorly. And unless you have selection at that level, basically, then you'll get monsters, as opposed to social organizations that actually function well in the context of their environment. So these are all real parallelisms between this biological level of organization and what takes place in humans.

AA: David, I have to say I have a certain fondness for monsters, though. Monsters of all kinds. I think oftentimes they do hold these keys to understanding what mechanisms are. Sometimes they have new solutions embedded within them, even though at first they look like they're something that's nonfunctional. And even in human organizations, sometimes, yeah, they're really messed up and weird, but sometimes you get really interesting things coming out of completely, seemingly dysfunctional ways that humans relate. So I'm not going to completely dismiss monsters, because they do have a special place in my heart. And it's not just because of the zombie thing, David, it's not just the zombie thing.

DSW: Yeah. Well, I mean, this is the second or third time we've come up against a problem which surfaces throughout this series, is that there's forms of superorganisms not worth wanting. There's forms of superorganisms that are downright scary. We don't want societies to treat us like skin cells, thank you very much. We don't want societies that basically make that kind of valuation so that if you're not the right type, then you're devalued, which I think so if you're a monster. Then for some reason, you're going to be discarded by that society. We want something, as did Teilhard, that respects individual freedom and dignity and autonomy.

And so there's real moral tensions here when we talk about group-level organisms, as to the kinds which exist, and the kinds which we want for ourselves, and among the dystopias that we could consider, first of all, collapses so that there's no functional organization at all, or things that work very well, thank you very much, but in a way which is just horrifying as far as our respect for individuals is concerned. And they all could happen. It's like threading the needle as to the kind of superorganism we want, or we think is worth wanting. So this is very fraught with tensions.

AA: Yeah. A civilization that was based on the principles of multicellularity would be a absolutely terrifying authoritarian government. So we don't want that.

DSW: But there's reasons. And this might... Interesting to see how this maybe translates back into cellular interactions, as we find our way towards discussing multicellularity and such things as cancer, is that there's a reason why human society needs to be egalitarian. And that reason is most evident in small groups because in the episodes that focus on human evolution, it's all about cooperation, and cooperation is all about social control, preventing bullying behaviors. The one danger in human social groups was the danger of being bullied, of being pushed around. And the only reason that small human groups cooperate is because they have the wherewithal to collectively gang up against bullies. And so that means, and this is as you know, it's called the fraternal style of major evolutionary transitions. Unless individuals are capable of standing up for themselves, then they will be bullied.

And so, this requires a kind of an egalitarianism which exists in most small-scale societies, and that egalitarianism needs to be retained as the scale of society increases. If it's not, then we get despotism. That's what we get. So the optimistic message, as I like to see it, is that thanks to the fact that we're not genetically highly related to the other people in our society, the only way for us to operate as a unit is with various forms of egalitarian social control. Like the rules of meiosis, as a biological example. Every gene has a fair chance of making it to the next generation, and if you don't have that, you have meiotic drive, and then the organism ceases to function.

AA: Even that is probably a result of a huge multi-layer regulation system that keeps meiotic drive from getting out of control, because that's such a weak point in the transmission. So I think that there are many phenomena in biology where it's just been assumed that oh, this problem is solved, but really underlying that are all of these layers, multiple layers and checks and balances of regulation that can break. And that also, if we don't appreciate that it works because we have those layers of regulation, then we just assume that the problem of cooperation is solved just by virtue of genetic relatedness, which I think is a problem that is, or that assumption is quite widespread in the evolution and behavior community, almost like, oh, you're genetically identical, so there's not a problem. It's like, well, no, you could still have all these issues of cheating and exploitation, unless there was selection that acted to create these regulatory mechanisms on this higher level.

DSW: And the whole vocabulary of human social control has been transported. We have parliaments of genes, we have policing, we have sheriff genes. And so the whole vocabulary of human collective action and social control is now part of the parlance of people who work on that. So, Michael, you were about to say something.

ML: Well, two quick things. One is that solving this problem by genetic relatedness is really interesting because on the one hand, we have situations where even in the same body, during embryonic development, all kinds of competition takes place, and in fact, it's necessary for coordination. Cells and tissues compete for informational signals, they compete for metabolic resources, and it's actually really critical to make sure that overall development goes correctly. And conversely, we make hybrids. We can make chimeras of cells of very distant origin, and they will cooperate.

Life is incredibly interoperable. The cells will cooperate with electronic devices, they will cooperate with other cells that belong to a completely different type of animal, then they will build something. We make frogolotls that are 50% frog, and let's say 50% axolotl. And they make something, the cells have no problem cooperating with each other, even though they're quite distinct genetics. And I think with this issue of what kind of societies we want and everything of that nature, it's important to realize how limited our view is in our ability to detect agency that is impacting us.

And what I mean by this is that we are very good at, as most animals are, at detecting agency outside of us in the three-dimensional worlds. All our senses point outwards, we look around and we say, okay, that is a rock, but that is a tiger, or prey or whatever it's going to be, and we can recognize that.

Imagine if we grew up with a sense that was biofeedback, if you knew at every given point what your pancreas was doing. I think we would have no trouble recognizing it as an intelligent agent solving problems, if we had access to the space that it works in, the actions that it takes. If we could see it working in physiological space, the way that we see animals running around in three-dimensional space, we would have a much better appreciation for what that intelligence looks like. And so that's sort of going downwards, but also going upwards. I think that it would be a Lovecraftian scale of horror if a single cell could realize that my god, I'm part of this organism, and the actions that I take in my space have been completely bent by... That space has been just distorted, in the relativity sense. That space has been distorted by the needs of this larger system of which-

AA: The Matrix, Michael. It's the Matrix.

ML: Right? I mean, it would be incredible. And of course, by definition, these parts don't have the wherewithal to understand what the goals are, to which their space is distorted. But I often wonder, are there any tools of information theory, of something that would allow you to at least have some idea that you're part of this larger scheme that is just distorting your options in some way? I think we're very bad at recognizing that, intrinsically.

AA: I mean, Michael, isn't that what we're all living in? A human society where we are shaped by the goals of others and the incentives that are around us. And our whole reality, especially if we live in a city and we go to a workplace where we're around other people, so much of our world is created by other humans, and that shapes what we can and can't do, and what the outcomes are for different courses of action that we might take.

ML: Yeah. That's absolutely true, but there's an even extra layer. Imagine for a moment that you were shrunk down to the size of a single cell, and you didn't know about embryonic development, you were placed inside of an early embryo. Frog, fish, human, doesn't matter. And you looked around, and you saw this incredible activity. Cells are dying, cells are randomly falling off, other stuff is happening, there's lots of noise. You could easily tell that your options were impacted by what your neighbors were doing, that there were signals. So you would for sure know that you are impacted, in exactly the way that you would say. What I think you wouldn't necessarily know, if you didn't already understand what embryonic development was, I don't think you'd have a clue that all of this giant rigmarole, this process with its noise and everything else, is going to reliably crank out a zebrafish every single time. I don't think you would catch that at all.

And I think we take this for granted when we say the acorns make oak trees, that's how it's supposed to be. But, we know that because we've seen it happen. We have zero capability right now to predict that. It's developmental biology. If you were shrunk down to a local perspective and seeing all the stuff that was going on, would you be able to tell that this is all part of making a zebrafish? I don't think you would. And so that extra step, going from knowing that yes, you are constrained by these million things, and that no doubt a lot of things are affecting what you're going to be able to do, to the next step of saying, my god, I'm part of something that actually has a goal state, it will reliably reach this, and have some ability to comprehend it to whatever extent our intellect allows. Right?

DSW: This speaks to the unintended component, the blind component of human cultural evolution, where cultures work without anyone knowing how they work. To some extent, and it's the intentional and the unintended, the blind and intentional components of human cultural evolution is the quintessential glass half-full. It's hard to know what to be more impressed by. The blind component, which leads to design that nobody designed, or the intentional component where people actually sit around, set about with their social arrangements and succeeded, to a degree, at accomplishing what they intended to build? There's excellent examples of both, but to focus on the blind component, then it results in exactly what you're saying, Michael. We're all participants in something that actually does result in order, to a degree, and we're completely unaware of it because we didn't design it. It's just what survived and what hung together, compared to the many things that fall apart.

So it's up to us as scientists to understand that, just the way with any other thing where instrumentation and theorizing is required to apprehend what we had no idea, that had no way of knowing, couldn't possibly have known if we didn't actually apprehend it scientifically. So before we can barrel on, I want to go back to the symbiotic cell theory, which has been covered in other episodes, but the idea that the nucleated cells, eukaryotic cells evolved as symbiotic communities of bacterial cells, thus documented for mitochondria and chloroplasts. But there's also a question as to whether any other aspects of eukaryotic cell structure, such as flagella or other organelles, also have separate origins. And I don't know, either one of you, if you have opinions on that, but just dwell a little bit on the symbiotic cell theory associated with Lynn Margulis, as we're talking at the cellular level.

AA: Well, I think that the evidence is quite convincing. And to me, one of the really fascinating pieces of it is this idea that because you built a eukaryotic cell out of these two elements, that where there was some level of conflict. There kind of still is. You could see in the way the genomes interact with each other, that there's some level of conflict. That you actually had a very efficient mechanism for cell death. One that you could turn on a dime, because I think it's hard to build into a cooperative system, this ability to self-destruct if things are getting out of hand. But by actually introducing genes where there was conflict, I think that might have actually been a really important piece of allowing for this cell death program, the apoptosis program to be built in at this most fundamental level.

So then you could actually build much larger bodies because you have this quick switch, like a kill switch for cells that are misbehaving. And I think it might have otherwise actually been too challenging to get into that kind of space evolutionarily, where you'd have a really solid kill switch built into each cell. You had to get something else in there that previously evolved to try to kill the cell. Put it in there under regulatory control. And then I think that might have actually allowed for much larger and more complex organization than was possible before.

ML: I think one of the upcoming fields that's going to have a lot of interesting things to say about that is bioengineering and chimerism. Because we are now putting things into cells that have never been there before, and so on. And both at the cell and tissue level, people make cyborgs and hybrots and all these different kinds of mixes of cellular material, and subcellular material, and designed engineered artifacts and so on. And what you see is this incredible ability to adapt and work with whatever's there. And I think it would surprise me not at all that early events in evolution would combine all these different things, because we see cells making use of stuff that is just incredibly alien to them. And no problem, they figure out how to get along with all these weird materials and so on.

And so there's this huge area of synthetic biology and bioengineering, where we're giving cells novel metabolic capabilities, novel signaling capabilities, novel computational capabilities, biomechanical capabilities, and all of this is highly interoperable. And so that plasticity of being able to work with whatever you've encountered, if you can make some sort of use out of it, I think I wouldn't be surprised at all.

DSW: Let's make cancer center stage here. And why is cancer a fundamental problem with multicellular organisms? Athena, why don't you go first, and then Michael.

AA: Sure. Yeah. So, cancer is a problem because multicellularity is fundamentally a cooperative system of cells. So in order to make this transition from unicellular life, or aggregates of cells that were just operating together, to a truly multicellular structure and method of reproducing, there had to be some mechanisms in there that kept cells from replicating and monopolizing the resources in a way that would limit the ability of that organism to function. Just like a cooperative group can be totally taken down by one exploiter, same thing with a group of cells. If there's one cell in there that's replicating out of control, or monopolizing all the resources and creating waste products, that can make it not viable to have a cellular group. So the evolution of multicellularity itself would not have really been possible without checks and balances on these potential ways that cells might cheat.

I mentioned replicating too much, monopolizing resources, creating lots of waste products, not participating in the creation and maintenance of the extracellular environment. That's actually a really important piece as well. And apoptosis also. So this controlled cell death, when the cell, if it's infected, or if it's mutated, or if just developmentally it makes sense for those cells to fall away. If there isn't a mechanism for apoptosis that's reliable, and that's going back to the earlier part of the conversation about how important that it is to have a really fail-proof apoptosis mechanism. Those are just necessary pieces to have multicellular function. And so really built into multicellularity is this vulnerability to cancer, but also all of these regulatory mechanisms that help to suppress it.

DSW: Yeah. Just, I wanted to emphasize one thing you said. On the one hand we have cheating, and on the other hand, we have just the coordination that's required to build an organism, even if everyone is willing. And you said, need to fall away, I think is one of the things that you said. That just in the process of building an organism, it's a need for there to be competition, as Michael has stressed, and for there to be cell death, even before we get to cheating, just because that's the way the organism is molded. So there's coordination, and then there's cheating. And a lot of what you were saying about such things as cell death and so on, would need to take place just for the coordination part of it, that cells have to die under some circumstances in order to build the organism. So Michael, why don't you take your turn on this?

ML: Yeah, so I'll start with the cheating portion and then tell our story of cancer and the lens through which I see this process. Let's go back to this idea of cells being joined by these gap junctions. One of the things that it does is, it greatly increases cooperation because when your internal environment is directly connected to the internal environment of your neighbor, there's no cheating possible, because anything your neighbor does is felt by you, anything you do is felt by your neighbor. It's like an effective sort of karma system, that there's no cheating because you are not distinct anymore. In many ways you have ceased to be, and now there's this new greater entity that you'd be cheating yourself in effect, if you poisoned it or something like that.

So I'll tell this story of how we see cancer. I completely agree with Athena, it's a fundamental aspect of being multicellular and having this need for cooperation. So one thing that happens when you've got cells that particularly join into these electrical networks, is that these electrical networks are free, from physics, are subject to spontaneous symmetry breaking, and the organization of patterns that serve as large-scale homeostatic states, where now your homeostatic state is a pattern. It's a spatial arrangement, not a local metabolic state. So this facilitates certain kinds of morphogenesis, because downstream they control gene expression and cell movement and all that.

So if you, in your mind, and we sort of have some tools for doing this, if you visualize the scale of the self that is formed, or you start out with a single cell and you'd say, okay, the scale in terms of a spatiotemporal cognitive radius, this light cone that limits what you can work towards... There's some metabolic goals, and you have a little bit of memory going back, a little bit of predictive capacity going forward. Once you're this network, and you have these larger goals that you can work on, we're going to make a liver, or we're going to make a heart or whatever, because that's what this network is able to do. What happens is that boundary of the self scales, so that what happens, it goes up, so that now everybody's sort of cooperating on this one project.

And when that process is screwed up, and of course, one of the first steps of conversion, when you start to convert to cancer, one of the first things that happens is gap junctions close. Your gap junctions close, you lose electrical connectivity. As soon as you've lost electrical connectivity, that boundary between self and outside environment shrinks. Because now that cell, as far as it's concerned, everything it sees around itself? That's just an external environment, as far as it's concerned, right? It can go back to being an amoeba because now the boundary between me and everybody else is very tiny. I would say that these cells are not any more selfish than anything else. It's just, their selves have shrunk. Everybody's selfish it's just before your self was kind of huge and everybody was working towards it. Now your self is tiny and now you're back to your goals from those ancient days where, what are my goals, right? Each amoeba wants to be two amoeba and you want to go where life is good and sort of you got conversion and metastasis. And so just to say that we've tested this. And so one thing you can do is you can kickstart basically full on, for example, full on metastatic melanoma, in a tadpole, no carcinogens, no DNA damage, no oncogenes, just by disrupting the electrical state of a certain set of cells. And they just these melanocytes completely convert and they go crazy and take over the animal.

You can go in the opposite direction and you can put in really nasty oncogenes, human KRAS mutations that make tumors in these animals. And if you at the same time, very carefully manage the electrical state using optogenetics or injected channels, right. You will prevent tumorogenesis, despite the fact that the mutation is still there, the oncogene is blazingly expressed. There can be no tumor because you've forced those cells into their electrical connection. So, that informational boundary, that computational surface between the boundary of that self and the rest of the world is what sort of shrinks and grows both during development and evolution and cancer.

DSW: Cannot help, but notice the similarity with the fragmentation of social identities in modern life. Athena over to you.

AA: So, I'm just super excited listening to you talk about this, Michael, because it implies that there's a whole other level of health that people don't think about. They talk about physical health and mental health and emotional health, but how about informational computational health, right? Having your, the cells of your body, having proper electrical signaling, being in the right expression state. I mean, I feel that is mechanistically a unique piece of health, and it's something that people aren't talking about as a piece of health in and of itself.

DSW: What you just described, Michael, is so evocative of a situation where, how people behave depends on their social identities. And if they have a large social identity, then they, that functions well. And if their social identity shrinks, then they just carry on what they're doing. It's just a different us. It's a different them. And of course, that could play havoc at larger scale. So, do you see that parallelism? I mean, it's so striking to me.

ML: I see it. I think that we have to be careful exactly what you said earlier, which is that there are larger scale structures that you can be part of that are morally abhorrent and not really what we want to aim for. So, the real question is, being part of a larger structure tends to raise your overall IQ. The question is what are going to be the goals of these large scale structures? And in general, this is one of the things that I talk about when we do synthetic biology, we make xenobots and people want to know about when they're going to be cleaning out your arteries and cleaning up the ocean, whatever, that's fine. There will be specific applications.

But the bigger issue to me, the much bigger issue is to get started working on a science, which we don't yet have of figuring out where the goals of collective systems come from and how do we manage them, because we're not good at seeing them in advance, we're going to be surrounded by swarm robotics and the internet of things and all these structures, including and current social structures, that literally have their own goals that they're going to be trying to implement. We don't know what those are, we don't know how to envision them. We don't know how to predict them, and we don't know how to manage them. And getting a handle on that I think is super critical because these are highly non-linear emergent, surprising processes. We are not good at seeing them coming.

AA: That resonates a lot. I mean, I think we have a tendency as humans to be here's the thing, this is the goal. I'm going to put my finger on it exists here or there. But when it comes to these emergent systems the goal does not exist in one place. You can't localize it in one entity. And I think we are really bad at reasoning about that. And that leads to some serious vulnerabilities for us as a society, especially as things get more and more complex, and there are more and more opportunities for these emergent kinds of goals to come about.

DSW: So, I want to raise the issue of evolutionary processes that give rise to other evolutionary processes, and evolution as being something that could be intragenerational. In other words, taking place within a single organism, in addition to in intergenerational. We have some well known examples. So, the adaptive component of the immune system is the best known example, the rapid evolution of antibodies. I do a lot of work thinking of our behavioral system as an intergenerational process, what BF Skinner called selection by consequences. But there's more, and I think the more mechanistic you get than the more Darwinian processes, intragenerational Darwinian processes exist. And so I just wondered if you could add to that list of basically Darwinian processes that take place within the organism, in addition to the adaptive component of the immune system and our behavioral flexibility.

AA: Cancer is one that comes to mind, you have cells that are evolving within the body to, they sort of develop their own fitness function, you could say.

DSW: That's one that's unregulated. So that's an unregulated intergenerational, yeah.

AA: It's unregulated, although there are some really interesting ideas now. So, if you look at there are certain kinds of mutations notch mutations, where you'll often see an expanse that takes up areas where there's sort of been injury. And me and Carlo Maley, we have a hypothesis that some of these clonal expansions in some places might actually be almost a plan B for if you have a wound and you need to have a mutation that will lead to rapid expansion of a cell clone that isn't likely to progress to cancer. So, that it doesn't get taken up by one that might be more out of control.

So, there are hypotheses out there that some aspects of somatic evolution might actually be under some control in order to crowd out those processes that are less under control. I think it's a really exciting venue for new work, but there's because there definitely are some clonal expansions that are much less likely to lead to cancer than others. So, it could very well be that there's some level of control there that flips some of those switches to just get the clonal expansion and migration without necessarily increasing overall cancer risk that much.

ML: I'm not sure I have any good examples of new phenomena like that. Although, I have one that I can tell you, that's kind of fun that sounds like that, but I actually think isn't at all. And I think that's important. What I do like about this idea of looking for evolution within the organism. One thing I'm really interested in this aspect of ontogeny recapitulating phylogeny in the journey across the Cartesian cut, right? We all start life as one cell. And people like to look at this one cell as a piece of physics and it's kind of this physical analog system. And then at some point it becomes, at least in some cases let's say human or similar, it becomes this novel creature that has cognition, first person perspective, and all that. We all got there in a very smooth continuous way, the same way that evolution got us there from single cells all the way to more complex forms.

And so I really like this, this idea of trying to understand what are the invariances in this journey, right? And there's, and the fact that in both cases, there is no, what I think is instructive in both cases is there's no one place where biology says, draw a sharp line here, everything before that was just physics, and after that, ah, now you have computation, representation, true preferences, true motivation, all that. There's nowhere that you can draw a sharp line, right? Neither in evolution nor in development. And so I'm interested in asking what aspects of that kind of scaling and what processes operate at the very large scale, and then at the very small scale to allow that kind of journey to take place?

DSW: So, I want to talk a little bit more about cancer defenses, the mechanisms that evolve to prevent cancer, which are extensive. And I guess they just shade into the immune system, which also protects against external infectious agents. And so let's talk about that, defenses basically, presumably they evolved by individual level selection. Organisms that had these defenses survived and reproduced better than those that didn't, and it leads to just fascinating conjecture, Athena, that I just love as part of your work. It’s in long lived organism, that cancer defensive mechanisms are needed the most. And so if we want to find species where cancer defense is exceptionally good, we should go to the whales and the elephants, not the mice. And, and so this kind of opens up a kind of across species comparison of cancer, which did not exist for cancer research for the most part, but now does. So let's talk about the elaborate defenses against cancer that have evolved. And of course the need to have something similar for large scale human society. So, speak to us a little bit about that, Athena.

AA: Well, one of the things to me that is really awesome about cancer defense mechanisms is that they exist on many different levels. So, you have basically cell intrinsic mechanisms, so things that are inside the cell encoded into the DNA of the cell that regulate the cells behavior. So, for example, the gene TP53, which is essentially like an information processing hub. So, Michael, to kind of connect to your work a little bit, I mean, this isn't electrical information, it's mostly gene products, but essentially what this hub is doing is listening in on all of these activities that are going on inside the cell. What's its metabolism like, what's going on with the cell cycle. What other proteins is it producing?

So, by listening in on everything, and then it’s essentially deciding, is this cell behaving as it should for being a part of a multicellular organism in this particular spot? And if the answer is no, then it initiates this DNA repair process. If that fails, then it initiates apoptosis, cellular self destruction. So, you have this internal process. It’s all intrinsic, almost the cellular conscience, you could think of it that way.

And then you have the neighborhood processes. So, all of the cells as Michael was talking about earlier, they're connected to the cells nearby them sometimes directly through these junctions where they're sharing proteins and gene products and signaling electrically. And sometimes just because they're near enough by that the factors that are being produced are in this sort of milieu and they're all affecting each other.

So, basically cells are constantly monitoring their neighbors. I mean, they're like the nosiest neighbors you've ever had. I mean, they're like in your house that nosy. And cells will not stay alive unless their neighbors are saying stay alive. So, they not only do they require not having apoptosis signals, right? So, if their neighbors say die, they die just like that. But, they're also listening, "do my neighbors want me alive?" And if their neighbors, they're not getting those signals from their neighbors, then they self destruct. So, they're super sensitive to their social environment, you could say, and all of the cells have these nosy neighbors. So, you have intrinsic, then you have these neighborhood level functions and then you have the systemic level. So, the immune system.

And so the immune cells are monitoring the whole body for regions where it seems like something is a little off, does it seem maybe there are cells that are replicating that shouldn't, is there some sort of chronic situation that's that just seems a little bit off. And if so then the immune cells will go to that region and try to probably, Michael, a lot of what they're doing is trying to re-regulate the signaling. I haven't really thought about it that way, but probably initially that's what they do. And then if that doesn't work, that's probably when they go into sort of the apoptosis mode. I wonder if people have looked for that actually looked at the immune system as an information, electrical signal regulating system in its first pass, because that would be way less costly than all the cellular destruction that we usually associate with what the immune system is doing.

ML: There is some work on that on macrophages and things of that nature. This is great. I love the story you just told. But the only thing I can add to that is my favorite cancer suppression story, which is that apparently one of the best ways to suppress cancer is to have a really strong idea of what you should be doing instead, a really strong morphogenetic cascade. And it's really interesting. If you have a view, which I think is kind of crumbling, this view of cancer as very local checkpoint genes and suppression mechanisms and so on. You would predict that animals that have access to a very large pool of undifferentiated plastic cells, basically regenerators that they should be very prone to cancer.

And in fact, it's been said that the reason that humans are poor regenerators, is in order to avoid the kind of oncogenic risk that you would get for having cells that are just happy to proliferate and migrate and things like that. So, that would be the prediction. right? And on the other hand, if you have this idea that cancer is an escape from a strong morphogenetic memory of the network, you would predict that good regenerators should have very low incidences of cancer. So, that's nice because it's a case where these two theories make opposite predictions and that's been checked and good regenerators in fact have very low incidences of cancer. Having lots of plastic proliferative cells is no problem. If you have the information structure that harnesses them towards an adaptive outcome.

And so my favorite example of that is planaria. So, these planarian flatworms, there's a couple things about them. They are incredibly regenerative. So, you can cut them into pieces. The record I think is 275 pieces, something like that. And Thomas Hunt Morgan, 1903 or something, cut them into a couple of hundred pieces. And every piece knows exactly what a correct worm looks. They regenerate with incredible fidelity. They make the worm. They don't get cancer. Now a couple of things, (A) they've solved the aging problems. So, there's no such thing as an old planarian, they don't age. And so they are basically immortal. And so the worms that I have in my lab are in direct physical continuity with worms that were here almost half a billion years ago, these are them.

Now what they do do in order to proliferate, many of them can do sexual reproduction, but they usually don't. What they do is they tear themselves in half and then each half regrows the other half, and now you've got two worms. Now think about what this means from the perspective of inheritance. They basically don't have a Weisman's barrier because any mutation in the body, they have somatic inheritance, any mutation that doesn't kill a stem cell is inherited and replicated throughout the body once they divide. And that cell has to make more progeny to fill in the rest of the body, right? So, now for several 100 million years, they've been accumulating all these mutations, their genomes are an incredible mess. We don't even have a proper assembly in japonica, because the cells are mixaploid.

Every cell has a different number of chromosomes. You don't even know what you're sequencing because every cell is just a mess, right? Massive genomic diversity. They are champion regenerators. Their anatomy is rock solid every single time exactly correct. Now think about what this means for our claim to understand what the relationship between genomes and anatomy is? We don't have a clue here really because the genome is basically shot to hell, it's all over the place, every cell has a different genome, if you were to say, how genetically related they are. I mean, I don't even know how you estimate that because they have all kinds of mosaicism within a single animal.

And yet their anatomy is the strongest of any other species, the best control. So, to me this tells us something very important. That having genomic differences, genomic errors, that's not really primary. If you really want to defend against cancer, you should have the same type of large scale instructive pattern memories, stored biochemically, bioelectrically, biomechanically, whatever it's going to be, really strong so that any cell that has even an inkling of trying to do something different is overcome with this connection, with this functional connection to a collective, that's trying to maintain organs, tissues and so on. So, that's my favorite suppression story.

AA: Michael, this seems like it has huge implications for thinking about resilience broadly, right? I mean, I think that there's a sort of bias that many of us have, thinking that, well if you want to have a system that is robust and can replicate order and all of that the system itself has to be highly ordered and structured and that you have to be able to see how that order gets created. It's almost like a bureaucratic mindset about being resilient, but it seems very often the systems that are most resilient are the ones that actually go through challenges and then have to change what they're doing, or add these layers, or add this other component. And often what you end up with is something that's a total kludge, but that might be much better at being resilient than something that looks it has a structure that you could be like, "Oh, I see how that would do this and this and this", and that would lead to what the expected outcome is.

DSW: For sure. So, let me give some human examples and we can segue to our final segment about what all this mean for the future of humanity and the internet age. But when you were talking about nosy neighbor, cells, Athena, I was reminded of Jainism, which is one of the strongest and oldest religions in India, and the Jains have ascetics. So, basically they wander from house to house and they are fed by the households, but they operate in that capacity as the nosy neighbor. Their food restrictions are so great that they can't accept food unless it's pure from their standpoint. And so they'll inspect the entire household before they'll accept a morsel of food. And they also sharply question the occupants of the house and so on to see that they're good upstanding Jains. And if they're not, and if they leave without food, then the whole community sees that. And so this is actually a kind of a moral policing that takes place.

AA: Seem a little immune system almost.

DSW: Very much so. And then also one of my favorite episodes in this series is with Anne Clin who is a senior member of the volunteer workforce of Wikipedia. And we have a whole episode on how Wikipedia works as well as it does. And it's all of these checks and balances, protections against cheating, coordination. I mean, it's so reminiscent of our conversation. In my conversation with Anne, I said, "Anne, this looks the immune system. Do you actually talk about it that way? Do you actually use the metaphor, the immune system over there at Wikipedia?" And she said, "oh yeah". So, I mean, it's needed, something like this is needed for any form of higher level functional organization, be it a cell, a multicellular organism, or a well-functioning human society such as Wikipedia. And so with that, let's talk about what does all this mean? What thinking have you done on modern human society? It's electronically connected. And so electronic communication, not only does it apply to nervous systems, but as we've seen in this conversation way before.

So, what are your thoughts on just the idea of human cooperation, which of course does take place at impressively, large scales that were unimaginable even a few 1000 years ago, even a few 100 years ago, the concept of global cooperation was beyond the imagination, I think before several centuries ago. Now it's something which not only does it make sense, but it's the only thing that makes sense. And many people at least can appreciate that the need for a whole earth ethic. That's what needs to work well and everything underneath it needs to be coordinated. So, what do you think of all this? How do you apply these ideas to what Teilhard would've called the Omega Point or some kind of the prospect for some kind of global consciousness? Is it possible? I think we can agree it's not inevitable. So, what are your thoughts on all of those things?

AA: I think that the big challenge is a scaling issue. So, when we go from small scale societies to larger scale societies, if we look at how small scale societies regulate cooperation and also how they sort of have resilience built into almost their cultural DNA. We see these systems that are basically, we call them need based transfer systems where people will help each other, if they have the resources to help and if someone is in need. So, many of them still have, this is mine and this is yours. These cows are mine and these are yours, but if things get really bad and you need something, I will help you and I will give you a gift. And I won't expect you to pay me back.

And those kinds of systems we see across lots of small-scale societies. From Masai in East Africa. they call it osotua, to Fijians, fish and horticulture in the Fijian Islands. They have a system called kerikeri, same thing, this need-based transfer system. Now, when you start trying to scale up a need-based transfer system, though, it can get more challenging. So I think the best example we have from small-scale societies is actually in Fiji where oftentimes a whole village will get hit by a cyclone. And what this means is that villages will actually have relationships with each other that are sort of dyadic village-village relationships of this need-based transfer norm. One village will help a whole other village if they're hit because within the village people can't help each other.

DSW: Athena, I want to nail down a point that you made earlier. There has to be a history of this kind of thing happening for there to be adaptations to address it. So whole village destruction like that, that had to have happened for this kind of arrangement to have evolved. In a constant environment, would that never happen? Why would you expect a solution to it? So these are all called for. I just wanted to make that point. Now continue, please.

AA: Yeah, absolutely. And the scale of the disasters and the challenges is super important here. You want to have systems where the risk is being shared at scales where there is variation in the likelihood of being hit, right. So it's like because cyclones are sort of often narrow things, if one village is hit very, very hard, a neighboring village probably won't be hit as hard. But if you have destructive events that are much larger in scale, then you have to think, "Okay, well, what is the appropriate scale than for these risk-sharing arrangements?" And I think as society has scaled up, there was never any sort of conscious effort towards taking these kinds of risk management system and scaling them up. Because a lot of the economic growth, which fuels itself, of course, is based on how you make the most efficient transactions, where you get the most benefit for the least cost.

And there are a lot of things that capitalism works great for, including short-term efficiency. But what it doesn't do well is create systems where risk is being effectively managed on long temporal and large spatial scales. So I see this as a the critical issue for us as a society today. How do we grapple with the risk that we're facing locally, regionally, and on larger scales and then create the kinds of risk-sharing systems that have served our species so well for so long and translate that into something that makes sense for our global society? And I have a feeling it'll probably be some hybrid between need-based transfer type principles and more account keeping based principles, just because account-keeping principles are so pervasive and they work for a lot of the things that we need to happen in our society.

But I think we need to think very explicitly about how do we import some of these systems. And there's a lot of things to draw from in terms of we already have a very developed insurance industry. I think insurance industries, they're also sort of grappling with how to deal with you increasing risk and the increasing sort of entwinement of risk that characterizes our global society. So I don't know. I have hope that there's a realization that there are many problems here that maybe we can solve with a solution that incorporates some of these risk-sharing principles.

DSW: That's great, Athena. Thank you. Michael, your turn, and then I'll take my turn.

ML: Well, two things I wanted to say. One is if we go back to this idea that one of the central things that define a given agent, whether it be human or something smaller, is that is the scale of their goals. Then what we find is that humans are in a very unique position. If you are, let's say a goldfish, and let's say you have a cognitive horizon of 30 minutes, right. So you don't have any goals that are longer than 30 minutes long, let’s say, just for the sake of argument. Then one of the things that’s true, and I think it's true for most organisms, is that your lifespan is actually longer than that horizon, meaning that all of your goals are, in fact, achievable during your lifetime.

Humans are perhaps uniquely in the position where your goals, temporally and spatially, could be enormous. And many of them are absolutely not going to be achievable during your lifetime. Because your goals are now bigger than the roughly a hundred years that we have. And I have a feeling that this gives rise to a lot of unique issues in the realm of psychology and psychiatry that sit at the root of this situation, where for the first time, you now have a set of stressors, and actually stress a whole other thing we could talk about. You now have a set of stressors about certain goals that are guaranteed to not be achievable because you know that you have this smaller horizon of action. So that I think puts a lot of pressure on our activities.

The other thing that I'll just mention is I was once in a discussion with someone, it was a podcast interview, and I got done talking about basically trying to deconstruct a lot of these binary categories because someone was saying human this and human that. And I said, look, first, evolution-wise, it's very smooth all the way back. I'm not even sure what exactly we're talking about. And then in terms of modern hybridization, so you could have a creature that's 85% human brain, 15% implanted technology, and you can have any percentage that you want, right. So these are not binary categories. And so after all of this and trying to explain all the varieties that you would have a hard time saying if this was a human or not, he said, "Okay. Well, then what is that anyway? What's a human anyway, right?"

And I think it's a really good question. And I think what we can't rely on are criteria based on what exactly you're made of. So what components are in there, or how you got here? What percentage? I think it's ridiculous to draw line and say, "If you've got at least 92% of things that were evolved in the human lineage and the rest are technology, well, then you're a human." I don't think that flies. So if we have to come up with a more generic, more useful kind of a definition for what a human is to begin with, I think going back to this, my sort of focus on goals and things you're concerned about. What are the things that you're actually actively trying to achieve?

I think that probably what a human needs to be is the definition of a minimal cognitive horizon that has a minimal radius of compassion outwards. So I think, and this is very compatible with ideas in Buddhist thought and in other traditions where the point isn't what you're made of. The point is, how much care can you muster about the other intelligences and the other sentient beings around you? So if I had to define what a human is, I wouldn't talk about brain structure or time since speciation or any of that stuff. I would be looking at the set of what does your cognitive structure allow you to have as goals that are large enough that comprise other sentient beings? And if it's big enough, if it's bigger than a certain minimal level, which I'm not going to be the one to set that level, obviously. But bigger than a particular radius than you are a fully-fledged human, no matter what you're made of or how you got here.

And then we can very comfortably talk about things like diminished capacity, where, okay, something happened, and you're not in fact able to care about the welfare of the standard number of things that a human should be. And so, in court, certain things will happen. And it also lets us think about what I think is inevitable, which is the eventual appearance of creatures that have greater than human cognitive horizons, that may have, in fact, greater moral responsibility because they actually can, they physically can care about more things than a standard-issue human can care about. And again, there are traditions that talk about that, right. Expansion of consciousness to the point where you are actively able to care about the wellbeing of a large number of sentient creatures, that a normal person can't physically muster in the linear range of what you can actively care about.

And so those are the kinds of things that I think about. How do we define humans in terms of both moral responsibility, and in fact responsibility to these various different kind of creatures, based on outward-facing goals and what it is that they can care about with respect to other sentient beings?

DSW: Wow. Okay. Well, okay. So I have a bit of wrap-up to do here. I want to make a connection with another episode with Jim Coan and Garriy Shteynberg. Back to small scale society, and what Jim Coan calls social baseline theory, which is, or which takes note that humans never lived alone. They always lived in the context of cooperative groups, even when those groups were waring with other groups. And as a result of that, genetic evolution has basically baked into our brains and bodies, social resources in addition to personal resources. That when our brains and bodies make their myriad trade-off decisions, most of which are beneath consciousness, they're always factoring in social resources in addition to personal resources.

And in that sense, we are members of colonies like an ant in a colony much more so than we have thought, especially during the age of individualism, which is the last 70 years. So there's a really interesting connection there. And then another point I want to make, a major point, is that there's two very different mindsets for thinking about being part of something larger than oneself. One has a very negative valence like we're all in the matrix, or we're all going to be skin cells, or we're all members of totalitarian regimes.

But then there's a mindset with a highly positive valence, which is basically the religious and spiritual mindset of being part of something larger than ourselves and wanting to be part of something larger than ourselves. And it's at that point that you're more joyful basically in becoming part of this. This describes so many people's wartime experiences where you're so bonded to your fighting unit that you'll willingly die for it. Where you can take a long-term goal that's longer than your lifetime. We're equipped to do that. It happens all the time. So I think isn't it interesting, and of course it all boils down to whether the individual perceives that being part of something larger than themselves as benign or not for themselves. And when they think it's benign, then it's something that they willingly do. And it's probably the optimal human experience is to be part of something larger than yourself.

And at a more mundane level, I think, and speaking of it in high spiritual terms. Back to my interview with Anne Clin on Wikipedia. Why is that so engaging for the hundreds and thousands of people that just love working on a volunteer basis for Wikipedia? Well, it's because it's a prosocial cause. Wikipedia is doing such wonderful work, if you join it. And you get to interact with like-minded people. It's just so very fulfilling. They love becoming part of something larger than their selves, and they'll spend all of their discretionary time doing it. So there is this positive face. Athena, you want to chip in on this?

AA: Yeah. I want to just emphasize that point. I think that we have a deep desire to contribute, to be a part of something. And many years ago, right, I worked on this idea of the walkaway strategy. You're in a group, and you leave if there isn't sufficient cooperation. Well, I think there's another piece that's very similar in terms of the mechanism, but it's kind of the opposite. You will leave a group if you feel like you don't have an opportunity to contribute or if your contribution isn't valued. And I think that is a really deep desire that we as humans have is to be a part of something where we are contributing to make it better, at least for most humans. And that I think is kind of often not really acknowledged as a deep and legitimate human drive, but I think it's there, and Wikipedia is such a great example of that.

DSW: And I think that when it comes to the scale of our social identities, the scale of our compassion, which Michael, you were talking about, there's such flexibility there because so many of our groups are socially constructed, almost all of them. Then the identity is that flexible. We could regard ourselves as citizens of our nations, or religions, or as small groups. Why not the whole earth? In fact, there's no more stretch to the imagination to consider ourselves, first and foremost, human beings and citizens of the world, than first and foremost, Democrats or Republicans are all social constructions. But then comes, of course, to what actually we do on that basis and how those things compete with other forms of behavior? At the end of the day, it comes down to a form of cultural selection, which needs to be itself culturally constructed.

So that half-full glass with the blind component and the intentional component of cultural evolution needs to become more intentional than ever before. And it needs to be experimental because we don't know what the consequences of our actions are. Really, all we can do is make informed guesses. Then we have to experiment again and again and again, and then we have to be mindful about what we're trying to accomplish and then selecting the things that get us in that direction.

So it ends up being essentially an intentional form of cultural evolution with the welfare of the whole system in mind and everything underneath that being appropriately coordinated. It's basically a managed process of cultural evolution at the global scale is what is needed, and nothing less will do is the way I frame it. So please share your thoughts. Let's have a wrap-up session. And then this has been such a wonderful addition to our series on the Science of The Noosphere. Athena, you want to go first and then Michael?

AA: Well, I think that this whole question of what do we need to do in terms of cultural infrastructure to help us be more resilient is really the question and the problem of our generation. And if we kind of just assume, well, the tech billionaires will figure out how to solve it with technology and money and becoming interplanetary, or whatever. Then I think we've kind of lost a lot of our autonomy to create the kinds of systems that we want that not only will help us be more resilient on a really large scale level, but also will help to feed us as humans that need to be embedded in social environments where we do feel like we have a safety net.

And I think it's a really fundamental thing that we have built into our evolutionary heritage. Because humans that were not embedded in social environments where they had safety nets were not as likely to make it as those who did, let alone societies that had cultural structures that had those safety nets. Those were much more likely to do well in the sort of cultural evolution stage.

So I think that thinking about how can we make that transition and be intentional about the structures that we're introducing on these regional, global levels and think about the scales, and have a structure where you have multiple levels that solve the problems that are at those different levels. Using some principles, kind of like some of these principles from Ostrom and some of these principles from small-scale societies that use need-based transfers, I think, could be really great.

And I'd love to see a more purposeful discussion about how do we implement, or at least test out some of these strategies on a larger scale, because we don't necessarily want to say, "Okay, here's the one solution. Now everybody just needs to do this." Right? You want to say, "Okay, here, well, we could try this out at a higher scale here with these partners and then see how it goes." So that's what I would like to see.

ML: Well, here's what I'd like to see as well. This is kind of my wish list. What I think we need of to optimize that relationship between our individual agency and the benefits of belonging to various large-scale structures is the ability to have some inkling into what the goals of the structure are so we can choose which ones we join. So I think a huge limitation on our agency is that we join things having absolutely no idea what our contribution to them is going to lead to. What the group is doing. It used to be easy.

You could look and you could say, "Okay, I could be part of these folks doing what they do or I could be part of these folks doing what they do." That it is now almost impossible to know all of the choices that you make in terms of what you consume online. What kind of products you interact with. What groups you hang around with. Extremely difficult to know what the end result of all of that is going to be.

So I look forward to a future where we have tools, initially, conceptual apparatus because this doesn't exist yet. So this needs to be developed and ultimately some sort of tool, whether that be some AI type of visualizer or something that enables you to make better informed choices about which kinds of things you want to be a part of. And this means not just kind of the obvious stuff.

I mean, all of us make choices about what groups we join. That's kind of obvious. But really developing a science of predicting the goals and the cognitive capacities of large scale kinds of systems that we are terrible at recognizing as large agents. We don't understand ways in which they're going to deform our action space. All of these things that are very hard to predict and completely not obvious.

So that's what I look forward to. I look forward to all of us sub-units having choice about what larger thing we're going to be a part of.

DSW: So this has been an amazing conversation, and a wonderful addition to our Science of The Noosphere series, and wonderful affirmation of the fact that when we tell this story, we need to include the entire arc from really the origin of life all the way up to the Internet Age. And this level that we've been dwelling on, of the single cell and multicellularity, is so informative. Who would've known. I think that almost everyone who listens to this is just going to be amazed at what we've been covering and its relevance to our current moment in human cultural evolution. So thank you so much, and I had a great time.

ML: Thank you so much.

AA: I had a great time too.

ML: I did too. Yeah. Great conversation. Thank you.