

## SCIENCE OF THE NOOSPHERE

**Francis Heylighen**

with

**David Sloan Wilson**

**David Sloan Wilson:** So Francis, I often reflect upon how much has happened in the history of ideas since the second half of the 20th century which couldn't have happened before. I think that's true for evolutionary biology, my specialty, and also for complexity which is your specialty, and which among other things require just the advent of widespread computing before we had the tools to even think about it or model complexity.

Because you go way back on this topic, I'd love for you to begin with just your own personal story as to how you became interested in complex systems, and against the background of really historic development in ideas. I think in some ways the enlightenment is still in progress. So, please just tell us your story against the background of these big ideas.

**Francis Heylighen:** Well, when I was in high school, actually I already started thinking about these things and I had even formulated what I called the generalized principle of natural selection which is that evolution to variation, selection is not limited to biological species but any system, physical, psychological, cultural, undergoes variation, selection.

So, I had this very broad train of interest but I didn't know very well what to study. I settled on studying physics basically for the reason that it was the most difficult discipline. So I thought, "If I can do this at the university, I can take all this on my own." That is more or less what happened. So I studied theoretical physics which is pretty hard but if afterwards, I was interested in psychology or in biology. I could do it on my own.

At the same time when I started studying physics, I did it with a kind of critical attitude in the sense that I was from the beginning skeptical about the traditional reductionist, deterministic, mechanistic culture. It's what I call the Newtonian worldview. At that moment in physics, there wasn't yet anything like complexity science. So, I started looking outside of physics.

What I discovered that was closest to the things I knew was systems theory and cybernetics. So in systems theory and cybernetics, there was an attempt to make a more mathematical theory using some methods from physics but without the assumption that everything is reducible to particles moving in space, the way the traditional physical worldview is.

It looked at systems and these systems could be biological systems, social systems, cultural systems that were connected by inputs and outputs, let's say systems theory point of view. And then cybernetics added to that, the idea that these systems could be goal-directed, that they were kind of like agencies. That they weren't just possibly receiving some input and transforming it into an output, that they were trying to reach certain goals. So that was a framework that to me sounded very natural, and which I recognized myself.

I also came into contact with a couple of colleagues working in that who actually in cybernetics were longer at that time than I was, and in particular came in contact with two guys, Valentin Turchin and Cliff Joslyn who had started something. It was in the very beginning stages, they called it the Principia Cybernetica Project.

So the idea is cybernetics and systems theory have all these great ideas. In principle, these great ideas allow us to unify scientific disciplines with the idea of general systems. It should be a theory that applies

to any kinds of systems. In practice, the domain is anything but unified. It's kind of a mishmash of all kind of different approaches, these different kinds of theories.

It was of course very ambitious, was a little bit like the Principia Mathematica of Russell and Whitehead where they tried to create a foundation for mathematics. So, we wanted to create a foundation for cybernetics. We called this as Principia Cybernetica Project, and I assume we made some quite good progress in the sense that we made a website where we formulated the basic concept, the basic principles that we organized a little bit like a semantic network.

It was one of the first serious websites in the world. It's 1993 actually. So, it had a lot of influence. A lot of people learned about these things from our website, but that's not yet what is conventionally known as complexity theory. So systems theory and cybernetics in principle are about complex systems but in practice, without computers or the models who could make mathematically. They had to be pretty simple. So, it's early in the 1980s that with the Santa Fe Institute, people started to model complex systems with computers. That made me also see that the traditional cybernetics and systems paradigm lacked something.

The traditional cybernetics and systems paradigm is kind of you have a well-defined system and you have a control mechanism on top of that system. On top of that control mechanism, you can have another control mechanism. So, you build hierarchies of systems within systems. All very fine, but the assumption is that the system has a clear boundary. That there is this one system, and that system has particular subsystems.

What we now saw in this new paradigm which is best called complex adaptive systems. I think the guy who was most responsible for that is John Holland, who was also associated with the Santa Fe Institute, is that systems consist of agents and the agents themselves are pretty simple. If you put lots of agents together, they have all these kind of very complex behavior.

There are self-organizations. There are positive feedbacks. There are negative feedbacks. There are evolutionary processes going on and that led to domains like artificial life where you try to simulate a whole ecosystem of evolving organisms. It was a very interesting approach to complexity that I would say complemented the one of cybernetics and general systems theory.

**DSW:** Let me cut in here. This is all great. I just wanted to drill down on a certain point before you continue, which is the fact that former analytical models are both enabling and constraining. On the one hand of course, they're so much more precise than verbal modeling. So they were huge events over verbal modeling, but they have to make so many simplifying assumptions that they end up being a de facto denial of complexity.

When I talk about this, the difference between the two-body problem and the three-body problem in physics, you hit this complexity wall beyond which formal analytical models can't go. So this embracing mathematics, formal mathematical models, for example in the whole field of economics, which is if you can't model it formally, then it doesn't have much of an impact. That's the sense in which computer simulation models were able to go beyond the complexity wall, and yet they were definitely frowned upon.

James Gleick in Chaos, I think made this point is that there is such a prestige and hubris associated with formal proofs that computer simulation models initially were looked down upon as a kind of a second class citizen. So finally now we've realized that there's no alternative to that. So, that's my interpretation. I just wondered if you agree or if you'd like to elaborate upon the account that I just gave.

**FH:** Well, I agree but I want to elaborate which is that the computer simulations themselves are also reductionistic. One of the reasons I said why I started to do physics but with a critical stance was because

I intuitively knew that these reductionistic, deterministic models, that they weren't right. They couldn't describe something like let's say living organisms or societies.

Even in physics itself, a large part of the history of physics in the 20th century is showing that all these deterministic assumptions which we thought we had at the end of the 19th century don't work. Quantum mechanics with the uncertainty principle, relativity theory saying that even space and time are not absolute. Chaos theory, of course showing that the slightest nonlinearity in your system and all kind of thing become, in practice, unpredictable like you said the three-body problem.

What physics basically has been doing in a way was to prove that it couldn't solve a lot of problems. That was what happened with other impossibility proofs like the famous Godel, or the halting problem in computer science. So the mathematical models themselves, they show that they were incomplete. So, computer simulation then was possibly an alternative. But for the computer simulations, you also have to make a model. At least you have to make a simplification.

Now, the great innovation for me of this domain of complex adaptive systems, I prefer to use that term rather than complexity science because in complexity science, different people attempt to look at the different aspects and some people will speak about chaos theory as complexity science where chaos theory is not really complexity science in my view.

So complex adaptive systems, the greatest insight for me was that we can model things by having these agents. Agents are kind of little bits of programs that are programmed with what is called the condition-action rule. In this condition, the agent does that. In this condition, it does something else. Now, you just throw a whole bunch of agents together. Each agent, perceives some condition, performs an action. That action changes the conditions not only for itself but for all the other agents. So you now get all this kind of direct and indirect interactions between the agents, and that leads you to all kinds of very interesting results.

I already gave the examples of ecosystems. In economics, of course where the agents that are people who buy and sell things depending on what one agent sells and other one may buy. The prices may increase and you may have all these kinds of effects on the stock exchange. That insight is a very interesting paradigm that illuminates a lot. Of course, it depends on what do you choose as your element. What are your agents and what do you program your agents to do? What people typically do is they have some kind of a simplistic view of what the system is.

They program their agents. They see that maybe the simulation doesn't do what they want, so they play a little with the rules until they get something that looks right. Then they say, "Well, we have proven that this and this will happen," but actually they have just proven that this simulation of these agents with these rules does that thing.

**DSW:** Let me take my term, Francis. I'm really enjoying this conversation. So, here is my turn. I'd like to spend more time on the concept of complex adaptive systems and actually distinguish a number of varieties. Let's begin with Conway's famous Game of Life. In this case, they're not actually not agents. They're just positions on a grid which can exist in an on or an off state. Then we put in simple rules as to for that on and off state basically based on neighboring cells, and their on and off state. So you put on those simple rules. Then out comes this amazing diversity of outcomes.

So, there's your point about basically agents with simple rules producing this amazing diversity of outcomes. But there's nothing adaptive about those rules. I mean, those aren't evolutionary rules. These agents aren't organisms in any sense of the word. So we can make the point about agents with simple rules of behavior and then that leads to all of these amazing emergent properties.

We should not use the word adaptive yet because there's nothing about those agents that's adaptive. That's just an arbitrary set of rules which are applied by the modeler.

**FH:** It depends. The Game of Life is of course the most deterministic kind of system. So a lot of simulations have been using so-called cellular automata of which the Game of Life is one variety. I am personally very skeptical about whether cellular automata can teach us much about the real world because they are so simple and deterministic, that really the only thing they prove is that even with simple deterministic goals, you can get very complicated things. That's was an important insight. Once you're there, you need to go further.

Once you start to speak about agents, you can have different types of agents. Of course, you can have adaptive agents. That means agent that for example undergo variation, selection. That's the classical work of John Holland. He has rules that are formalized as things of ones and zeroes. Then he randomly changes some of these ones and zeroes, and he has a selection criterion. Then, the agents with the right rules may survive and the other ones get eliminated. That's one way to make them adaptive.

Actually, you can have adaptive processes even when your agents themselves are not adaptive. That for me personally, that's the most inspiring idea. You put a number of agents in an environment. The agents can change the state of the environment. They can do things. For example, suppose you have an environment and there are different sources of food in there and an agent can either eat a piece of food, or transform it into a different piece of food, or produce a piece of food.

The adaptivity is in effect that each agent has its own rules deterministically, but it doesn't know what the other agents are doing. The other agents have their own deterministic rules. Because they are in the same environment, what the one does affects the others but you don't know a priori, how that will happen. I think the novelty comes from the fact that you have independent agents, each of which is deterministic but their interaction is not deterministic. The interaction has to adapt to what all the other agents are doing.

So, it's a kind of a model of coevolution. Each agent does whatever it's programmed to do, but to do that, it has to take into account what the other agents do. And that changes all the time. So, you actually get the very adaptive systems. That is why I think the term complex adaptive system is not badly chosen.

**DSW:** There's two meanings. One meaning is, is that the agents within the system are following adaptive strategies. The other meaning is, is that the system as a whole is adaptive in some sense. So, I'd like to get your opinion on that distinction. If we're going to use the word adaptive, basically adaptation can exist at two levels. The agents within the system or the system as a whole.

In any practical sense, we're trying to create systems that are adaptive as whole systems. That's any kind of policy objective is to create whole systems that are adaptive— economic systems, environmental systems, social systems. So, how do we get from systems composed of agents following adaptive strategies, their own adaptive strategies to systems that are adaptive as whole systems?

**FH:** As I said you don't even need to have adaptive agents to have an adaptive system because the adaptive system is the whole of all the agents evolving and mutually adapting. One strategy that helped this whole adaptation, and that is I think one of the works that I'm most proud of. I have written several papers about it. I didn't invent it, but it is the concept of stigmergy.

What is stigmergy? Stigmergy means that an agent does something that leaves a trace in the environment. That environment is shared with the other agents. These other agents can, if you wish, read the traces left by previous agents, and then they will build on that. They will react on that. In another conversation, we spoke about Wikipedia working according to this principle. The different agents here are people who each write a little bit of texts in Wikipedia. These traces remain in Wikipedia. They are public, and they incite other people to add something to that.

The same applies at all of these levels even with very simple stupid agents and the classic example are the ants that leave pheromone traces. These pheromone traces then tell other ants what to do. The

even simpler case that was the one that gave the name of stigmergy are the termite hills. Termites, when they build a cathedral-like termite hill, they just tell by dropping a little bit of mud here or there.

The mud contains pheromones which attracts other termites and thus termites tend to add their mud where there is already the most termites. So, you get a positive feedback. The higher a part with mud is, the more termites it attracts. It's again, this principle of stigmergy. The trace left by the activity of one agent, one termite, incites other agents to add to that. Adding in this case is pretty obvious, it's just doing more of the same.

In the case of Wikipedia, the adding means changing, correcting, whatever it is, the trace in a shared environment, now provides a kind of a template that is constantly being changed by the other agents like that. We have a self-organizing system that can be very adaptive and very coordinated even though the agent themselves are stupid.

**DSW:** If we were to understand the mechanisms of development of a single organism in enough detail, we find something just like that. This brings up the concept of the organism is really a central concept and metaphor because with an organism, we have something which is manifestly adaptive as a whole system, in which everything that makes it up is of course contributing to that adaptive whole but in a way that is very distributed. As you've written, even the brain is very distributed. There's no homunculus within the brain. It's all distributed.

What we're working towards, especially with The Human Energy Project is the concept of the noosphere and ultimately the whole earth is some kind of superorganism. Before we can talk about superorganisms, we really should get our idea straight about organisms. That's our master kind of a concept. So let's talk about organisms.

What you just said, you said this whole system can be adaptive without the agents being adaptive. Is that true for an organism? Can you make that statement for an organism? In what sense are the elements that make up an organism, the cells, or the organs, are not adaptive in their contribution to the functioning of the whole organism?

**FH:** Of course in an organism, the way we know it by the logic of organisms, they are a result of many different evolutionary transitions where each time you started with a system that had some degree of adaptivity. So, I'm not going to claim that the multicellular organism is only adaptive at the multicellular level. The cells themselves obviously are also adaptive.

Let's make a caricature or simplification. Let's assume that each cell is purely programmed by its DNA. So it has a DNA program that tells it, "If these molecules enter the cell, then you produce these other molecules to deal with it." In principle, that should be sufficient for the whole of this cell to coordinate. In my view, the critical term here is coordination. Each agent is capable of adaptive action. The agent itself does not need to be adaptive in the sense, they changes its rules. The agent is adaptive in the sense that if something changes in the environment, the agent will perform some action that makes the action that the agent appropriates to the environment.

The problem is if you have lots of agents, that each perform his actions, how do you coordinate them? I distinguished two forms of coordination. The one that I already described, that's the stigmergy one. The stigmergy one is you have this common medium. You drop signals in that medium. Everybody can read them, and whoever wants to react to what you have put in the medium can do so.

The equivalent in the multicellular organism would be hormones. A cell has some kind of a problem. It is programmed in this form of stress to release a certain type of hormone. That hormone goes to the bloodstream. Other cells that are specialized in reading that kind of hormone now will react by maybe releasing some other chemicals or maybe they will change some of the activities. Maybe your heart rate

will go up. Or you will start sweating, all because of this one hormone that was deposited in the bloodstream. That's one way of coordinating the activities.

The other way, and that's the most sophisticated way is the one we find in the brain. In the brain, one neuron if it has let's say some kind of a problem, instead of depositing just the neurotransmitter that everybody can read, it will send a signal to particular other neurons that is connected to it. Then you get the network, and in the network, of course you wanted the right signals, go to the right agents.

Then you get this difficult problem that people in your network in theory are trying to solve. How do you get that network to self-organize so that it becomes effective in the sense that all the neurons together collectively are solving these difficult problems that single neurons cannot solve? So we know a number of algorithms that do that of which the most basic one is reinforcing whichever connection works, meaning that it did something that you wanted.

Let's say these are roughly the two prototypes of coordination. Either stigmergy, you just leave a message in a medium that everybody can read, or neural network, you send a particular message to one or more particular agents. If that's the right agents, the connection is reinforced. The next time, they are more likely to get this message.

**DSW:** Let's bring this to the noosphere, the concept of the noosphere. Maybe you could just define the noosphere for us and also update Teilhard. There was his vision, and then there's where we are today. It doesn't have to be exactly the same. Nobody is clairvoyant. So what was Teilhard's vision of the noosphere, and then how would you describe it in modern terms, noting both the similarities and the differences?

**FH:** Maybe I can give a little bit more comment about my own history in how I got to these concepts. In Belgium, let's say the ideas of Teilhard have never been far apart. It's not that it was a mainstream philosophy but I heard about him. I never studied him really but it was in the background. By the way, the books of Teilhard were published by a Belgian, Maxfield, 12 years after his death. So, there was a certain tradition of maybe taking Teilhard seriously.

Then, I said that I came into contact with this cybernetician, Valentin Turchin, who was a Russian that emigrated to the US. He had written a book that was called the Phenomenon of Science, obviously inspired by Teilhard's the Phenomenon of Man, which is not the best translation since the original French, *Le Phénomène Humain* would rather be *The Human Phenomena*.

Turchin was inspired by this and he was inspired particularly by this idea of the superorganism which he called the super-being. The idea was that you have these different cybernetic levels of higher level control, higher level intelligence. As you go through the different transitions, you come to the human level. The human level he described as the ability to think, meaning that you can reflect about things that you are not immediately confronted with it. You have not immediate experiences with. Animals can only react to the things they experience.

Then he felt a need to say yes, but that thinking is there are some kind of a control level beyond, another control level. Well, it had to be something like a noosphere. He didn't call it like it. We had different names for it. We called it the super-being or the superorganism, and that was about the time the world wide web was appearing. In our little group, *Principia Cybernetica*, me, Cliff Joslyn, Valentin Turchin, we were already anticipating that we would use the internet to communicate. We were looking in particular at some kind of a network-like hypertext-like system which didn't exist at the time in 1991.

Then I discovered almost by chance that Tim Berners-Lee had this great idea that just implemented what we wanted. Of course, we were even planning to maybe make some prototypes ourselves. So suddenly there is this world wide web and very nice, this world wide web. It has a structure that is kind of brain-like because what do you have in the brain, you have neurons connected by synapses. If you look at it at

a higher level, you have concepts that are collected by associations. That was also part of the inspiration of Tim Berners-Lee, hyperlinks are associations between texts. So, the web is a little bit kind of like a nervous system.

So what I did was, I immediately collected to the idea of Valentin Turchin of the super-being or the superorganism and came to the idea, well, actually, this world wide web may turn into what's eventually called the global brain. So there, you get this very nice idea of some kind of a super human structure that interconnects all the people that has a brain-like structure, but it actually also is not just the brain, there's also an anatomy, a physiology in there.

So the superorganism idea then became more concrete and like several people have done, one of the obvious inspirations was the Living Systems Theory of James Grier Miller, who had this notion of all living systems by which he meant not only biological systems but also social systems have these different kinds of sub-systems or critical functions. Critical functions like digestion, storage, memory, transport, distribution, et cetera, and these have obvious analogs in society.

The superorganism idea became pretty clear to me. The global brain idea was then the information processing part of the superorganism. At that moment, I wasn't using the concept of noosphere because the global brain seemed to be enough, and it seemed to be actually a more powerful metaphor maybe than the noosphere because the noosphere is a sphere of thought but it doesn't say how it functions.

Then later, there was the work I had been doing with Shima. I came to the conclusion that within this global brain, there are kind of two levels. There is you might say the anatomical level that all the different computers that are connected via links and information is sent from the one to the other, very brain-like. Then, there are also ideas that circulate.

The way I interpreted the noosphere, but I know that's only part of what is meant, it's more this space in which ideas circulate. That gets me more to the stigmergic paradigm than to the neural network paradigm. So the neural network paradigm is one thing is sent from A to B, and from B to C, and it's the right thing that needs to be sent from the right agent to the right agent. It's like I send an email to David to forward it's maybe to Terry. I don't want that mail to be read by anybody. It's targeted.

If we now look at things like Wikipedia or social media things, they are no longer targeted. You post something publicly and people can see it, and they can react on it, or they can ignore it, or they can publish it further. There you have much more of this kind of stigmergic type of organization where the particular linking structure doesn't matter that much. In that case, it's a very different dynamic. The dynamic that is better in some respects, worse in others.

**DSW:** Let me try to play that back and expand upon it, Francis. What you said I think is that if we look at for example non-human organisms and superorganisms, you see two different kinds of organization at play, stigmergy and neural network. So, they both contribute to the organism functioning as a whole. Now, if you look at the human case, we should see the same thing.

We should see something as brain-like and we should see that's stigmergic-like. They both have the effect of causing the whole system to function as a whole. Did I understand that correctly?

**FH:** Yeah. The both help the system to coordinate activities. They're both the kind of a communication medium through which activities can become more synergetic.

**DSW:** So, now I would like to make a new point, which is that although the internet and the internet age you might say is the current chapter of this, these ideas remain just as important as we go back in history. We look at such things as all the major events in history, roads, technology, institutions, bureaucracies. If you look at the cultural evolution of societies at progressively larger scales, then you'll find the equivalence. You'll find of both stigmergy and nervous system processes, systems of regulation, and so on. So, the concept of a superorganism or a noosphere exists at intermediate scales.

Actually, we might say, I do certainly. It does not yet exist at the global scale. We want it to. It might be on its way, but I would think in most respects, except in some very special cases such as the international space station or global efforts at solving the pandemic problem which are very feeble. That cooperation and coordination does not exist at the global scale, but it does exist at various intermediate scales.

If you look at the social systems that work the best, the best functioning nations, the best functioning corporations, the best functioning religions, there you will see some good examples of brain-like processes and stigmergic-like kinds of processes, at intermediate scales for the most part not yet at the global scale. That's what we need to create. What are your thoughts on that?

**FH:** I agree. What the internet has done is made these things in a sense much more visible because they happen so fast and because we have some ideas of the algorithms and the linking structures. It's easier to see. Actually as you set any well-organized social system whether it's a government, or an army, or a firm, has this internal channels of communications that are brain-like and there are quite a number of others who have been making that analogy.

For example, Stafford Beer, one of the founders of management cybernetics speaks about the brain of the firm or the firm as a brain. Herbert Spencer who was an evolutionary thinker and a father of sociology was looking at society as a superorganism, though he noted at that moment that there wasn't yet the equivalent of a brain there because he couldn't quite imagine something more world wide web like. So, lots of people had been making that analogy and that analogy is indeed correct.

At the global level on the other hand, I think maybe you are too pessimistic in the sense that when we hear about the things that go on, on the global level, there is what I call bad news bias. That is, what is reported in the media, all the things that go on, the wars, the terrorist attacks, the hurricanes, the pandemics. Then each time there is a tendency to blame like this hurricane wasn't dealt well with because they had saved some money on maybe on dikes and protections. What people don't see is all the problems that do get solved locally or internationally.

In terms of global coordination, I think the best example is science. It's not just the pandemic at this moment but the whole of science already since at least half a century is fully global, fully international. There isn't something like a Chinese science, and a Russian science, and an American science. It's just science. So I think there is a lot of coordination happening, but when the coordination functions the way you want it to, nobody notices it. I would be inclined to say we tend to focus on all the things that don't go well like attempts to detect our global warming, or the problem of the Taliban in Afghanistan, but all the other things.

The United Nations have this human development reports. Each type has some objective measures of how things progress. In each year, practically each thing has progressed. People have lived longer. People have become richer. People have gotten better educated. A lot of that is because of international aid in the poorest countries or simply because of the economic system which is also a coordination mechanism.

We all know the shortcomings of the market mechanism, but the market mechanism, the invisible hand is one of these self-organizing coordination mechanisms that is highly distributed and is highly globalized and can do quite a number of impressive things.

**DSW:** Well, here's my take on that, Francis, which enables me to be both optimistic and pessimistic at the same time. It comes right back to this distinction between complex system that's adaptive as a system as opposed to a complex system composed of agents following their respective adaptive strategies. Let me state it in two steps.

One step I think is so obviously that everyone I tell just nods their head because it's so obviously true. Namely, every systems engineer knows that you do not optimize a complex system by separately

optimizing its parts. No. You have to have the functioning of the whole complex system in mind and then you have to then for the most part experiment in order to get the parts of the system to coordinate.

So the systems engineering rule, you cannot optimize a complex system by separately optimizing the parts, because that ignores the interactions. Is that something that you can easily agree with? I would hope so.

**FH:** Well, I agree with the idea that you cannot optimize the whole by optimizing the parts. Actually, one of the things that I remember from our Principia Cybernetica Project was we made dictionary of cybernetics and systems which was not so much our work but we collected different definitions. One of them is called the principle of sub-optimization. It's not very well-known but it is exactly that. Optimizing a subsystem in general does not optimize the whole system.

I want to take issue with what you seem to imply which is that the engineer needs to optimize the whole. That means that the engineer would have some kind of what is called the utility function for the whole and would start from that. The problem is that the whole is usually much too complex to optimize in this way. So you need mechanisms of self-organization. For me, the most effective mechanism of self-organization is still local.

So the idea of self-organization, the definition of self-organization that is mostly commonly used is global order from local interactions. So the global, all that means you get some kind of a coordination at the global level. It's not necessarily an optimization because you can't really optimize anything that is non-trivial. It is the creation of a regime that functions pretty well at a global level, from local interaction and local interaction is this idea of the different agents mutually adapting to each other.

That's also one of the lessons of the complex adaptive systems that no agent knows the whole system but the agent knows its local environment. It knows that if it performs certain actions in that local environment, it will get in trouble with other agents, and therefore it experiments until it finds a way that doesn't create trouble in its local environment.

The same applies to all the other agents in the environment. They're all trying to find a way of dealing that doesn't get them in trouble with their local environment. If you take the local environments of all these agents, you have the whole environment because you go from agent to agent. Then what's my neighbor is for you, your second neighbor, and is for another one, the third neighbor. In the end, if all the neighbors adapt to each other, the whole system is adapted.

Now, I don't think that this mechanism will always avoid the global problems but it's a mechanism that is most commonly used in nature's local adaptation that then propagates to the global level.

**DSW:** Yes. Again, I agree with you, Francis. We're still on the same page. The idea that we can't just optimize the complex system in some top-down controlling engineering sense. The way I would put the way to do it, would be the constant experimentation because the whole system is so complex. Basically, you have to try out experimenting with the component interactions, but then you always have to be assessing it at the whole system level. So, there's some process in which your selection of the lower level processes has to be based on its effects on the whole system. There has to be some system level criteria for what you select or for what gets selected at the lower levels.

That's in fact, what takes place in a practical sense with any kind of complex systems engineering. You're always modeling the system or you're actually experimenting with the system. What you end up adopting or failing to adopt is based at the consequences of the whole system level. How can it be otherwise? So, system level adaptation requires system level selection. Of course, that involves a variation selection at the level of the components of the system. Would you agree with that characterization?

**FH:** I am not following you completely there. I'm not sure that it is always possible or desirable to compel local adaptations to the effect on the global level because in most systems, in nature, it's very difficult to know what's the effect at the global level. So, I do think there is a lot of power in these local adaptations. The idea of the invisible hand is very beautiful idea.

We know that it has its shortcoming and I can tell you if want precisely where these shortcomings come from, but the idea of the invisible hand is somewhere there's a lack of a certain thing. Supply is not sufficient for the demands, then those who can supply the thing that is lacking will produce more of it because they know that if there is more demand than supply, they will get a higher price. So, they are motivated to produce more of the thing that's lacking. They are not optimizing at the global level. They are just saying wherever there is a demand, they will try to provide the supply.

So, I think most processes work like that. There is a local demand for something and you produce something that provides that local demand, and all these local demands mutually reinforce each other. So, there will be cases like in the case of global warming where the system doesn't work that well, but in many cases, I think it works on this really local level.

**DSW:** Francis, I think we're actually zoning in on an important difference of opinion between us, which is a good thing because I think we now having identified it, we can maybe come to a resolution on it, if not in this conversation then over the longer term. First, let me speak as a biologist which is my home territory, then state the equivalent for a human system.

I think that in nature, the concept of an organism adaptation or superorganism is actually quite limited. We know that organisms are organisms. We know that beehives are organisms. We know that some ecosystems are organisms such as microbiomes. Very quickly, you enter a region in terms of such things as ecological communities or even social systems that I do not qualify as organisms or superorganisms.

For example, you get primate societies that are just plain despotic in human terms. The bad guys won and they just bully everyone, and that's the way it is. Okay? Or an ecosystem that's a basin of attraction. It's a configuration of species that's stable. It does not work well at the ecosystem level, just stable.

So, a lot of nature is like that. It does not deserve the term organism. Nothing adaptive is taking place except the war of all against all in some sense. Suffering is everywhere. The Buddha was right. Life is suffering and suffering is caused by craving and desire. So to be able to see the disorder in nature, in addition to the order of nature, is something that's very, very important.

Now when we move to human life, I think we see much the same thing. We see definitely organization, and that organization does extend above the individual level for sure. So, we do have societal superorganisms but we pass a point in which at threshold, in which what we see does not yet deserve the term superorganism or adaptation at that level.

We see agents conflicting and competing with each other. We see despotism. We see basically a world of suffering caused by craving and desires, colliding with each other. There's no invisible hand to save the day there. No invisible hand. The invisible hand is profoundly untrue, that agents pursuing their separate adaptive strategies somehow miraculously function well as a whole system.

My claim is that system level selection is required just as with systems level engineering. And that you get that, then the invisible hand does apply. That when a system has been selected as a system, it is indeed the case that the elements within that system do not have to have the welfare of the whole system in mind. The whole thing miraculously comes together as if led by an invisible hand, but that invisible hand was a process of system level selection. That's my position.

So we're not going to resolve it during this meeting, but I'm really happy to have clarified it. And if we disagree, then I look forward to continuing the conversation.

**FH:** Actually by you stating it likely that, I thought maybe of a solution. I wanted to introduce the research we are doing now that is funded by the Templeton Foundation which is called The Origins of Goal-Directedness. Because I just had an idea how this approach in this project may bridge the gap between what you say, whole system selection, and local selection. Which is the starting point of the project is that if you have a number of reactions and reactions are you might say simple agents, very simple agents. It's where an agent is something that if this is a case, then you do this action. If this is the case, then you do this action.

If you modeled it like chemical reactions but they don't need to be chemical. It can be anything that gives you and A plus B, gives a C plus D. It turns out that if you throw enough of these reactions together. They tend to form a so-called chemical organization which is a self-maintaining whole. Self-maintaining means that anything that is consumed is produced again.

That leads us to the important concept of autopoiesis. For me, of all the definitions of life or organism, I think the autopoiesis one is still the best one. An autopoietic system, is a system that produces itself, and in this way closes itself off to some degree from the environment. It creates its own boundary by having its internal logic of reactions, producing reactions, producing reactions that close in on themselves. So, it's kind of a generalized cycle. The A's produces B, that in term produces A's, it's a cycle. It's the thing producing itself. So, that is one way to define an organism.

Now, the closure—and now we come to the tricky part—the closure is what distinguishes the inside from the organism from the outside. That means everything happening in the environment. The closure is not a thermodynamical closure. Obviously, to have a living organism, some food must get in and some waste must get out. Then the question is, where are you going to draw the boundary?

Now in chemical organization theory, which is a mathematical formalism, you can define that mathematically. You can say, "This is an organization. This is the boundary. These are the things that are outside it. These are the things that are inside it." But organizations can be contained in other organizations. So, you can go to a bigger organization to a smaller organization. And you can go, let's say first from a smaller organization to a bigger one that encompasses it, and then from the bigger one again to another smaller one, which is not the same as the one before. So in principle, you can have an almost continuous evolution step by step, where you go from one organization that contains certain components, to another one that contains almost the same components but not quite. To another one which contains almost the same components but not quite. If you do this for several steps, you have process of evolution that may lead you to something that's very different from what you started out with.

But notice that this is a process of evolution in which nothing gets killed. So the traditional view in biology of natural selection is you need to generate a number of organisms. You kill off all the ones that are not adapted, and you keep some. The view in chemical organization theory or I might say more generally in self-organization is you create an organization. You perturb it in such a way that it can't go on in the same way. Instead of killing it, which means erasing it completely and starting anew, you turn it into a different one that is almost the same but not quite. If that new one is better adapted, fine. But if it turns out that this new one is still not very stable, so a new perturbation comes. Again, it has to choose.

**DSW:** You killed it. You killed the old one.

**FH:** You didn't really kill it.

**DSW:** There is death there.

**FH:** No, not really. It's more like the difference between societies. If in Belgium, a new government is elected to different parties, you could say, "Well, the system has changed but nobody will say that Belgium has died and now it's a new Belgium."

**DSW:** Francis, we might be splitting hairs...

**FH:** What I mean, it's selective but it's not killing off the whole thing. You select by removing the regime that didn't work, replacing it by a different regime, but a different regime is not a blank slate. You don't start from nothing.

**DSW:** Well, but that's not true in nature.

**FH:** You start from something that has most of the same ingredients but slightly adjusted.

**DSW:** So Francis, we might be splitting hairs at this point and I want to do a check with Alan as to whether there are important points. I think this has been a great conversation. Now just to take business, I'm really impressed. I have a lot of fun reading the business and management literature. I don't know if you read that literature, Francis.

One of the things I've discovered is that most innovation in the business world takes place not by companies changing as companies, but by them going under and being replaced by a new startup companies. So, there's real death there. I mean, legal corporations just ceased to exist and they're replaced by other corporations that were born in some sense and they're different.

Of course, they incorporated elements of other cultures and so on. If that's not a process of death, I don't know what would be. It's legal, social entities, no longer existing. So, I think that there's a sense in which cultural evolution of course requires the replacement of worse practices with better practices. And those worse practices, we hope, cease to exist. So, I'm not sure that there's a fundamental distinction to be made with the basic process of selection, cultural or genetic, being a process of differential survival and reproduction.

#### ***Break in conversation***

**DSW:** So Francis, we've been talking at a very high theoretical and scientific plane but we're part of a project called The Human Energy Project, which is very much based on Teilhard, and what we call a third story, that there's something about this which can be understood by many and really captivating to many. So there's a narrative end of this, which is very much centered on Teilhard per se. You described Teilhard as someone that you encountered early on and was kind of part of the air over there in Belgium. Then, you re-encountered and so on.

He wasn't really center stage, especially not in a narrative sense along with the concept of the noosphere. So, let's come back or let's finish up actually with the Human Energy Project, and the idea that Teilhard and the noosphere can become part of a third story that could be widely understood and valued by everyone. So, let's finish up that way.

**FH:** Well, the third story is also something that I have been working on without calling it that for ages. The research center where I work is called Center Leo Apostel. Leo Apostel was the famous Belgian philosopher who had kind of similar aims as Teilhard. I think he probably even was inspired to some degree though his background was different. He wanted what he called an integrated worldview.

The idea is science is fragmented into disciplines and sub-disciplines, that all the worldview from the regions don't give answers anymore to our big questions. What we need is a world a worldview that gives meaning to life that tells us what our position is here in the universe. Leo Apostel defined the concept of a worldview with six components. I'm not going to give them here, but in essence are about creating meaningful life.

Creating meaningful life, why do we need a third story? It's because the second story of science, it's not meaningful in the sense that it tells you that the universe just follows some laws that just happened to be there. We don't know why, and it's all like a clockwork. It's just running and you don't have anything

to do with it. It doesn't go anywhere particular, or in the newer version of the scientific worldview, it's all the end of an unpredictable and chaotic. There is no sense of direction.

What the third story should do is give us the sense of direction. The sense of direction, we don't want to go back to the first story where the first story is the traditional religious story where God has some kind of purpose for the universe and the story is that we are just fulfilling God's design. We don't want to go to that either. We want to go to a worldview where there's both a sense of direction but there's no predetermination.

If you look at the more than scientific world, if for example this worldview of evolutionary theory. The worldview of complex adaptive systems, the worldview of self-organization, you see that processes have directions but they are not deterministic. We cannot predict what evolution will give but we can predict that some things will not happen, and some things will end badly, and some things may work also.

For me, the third story is this idea that there is a directionality to evolution and Teilhard formulated in a pretty simple way that in this law of complexity consciousness says that during evolution, complexity increases. Something which I think most people would intuitively agree with, although some evolutionary biologists are skeptical about that.

The second one is even more interesting. It is that consciousness will increase. Then of course, you need to define what is this consciousness. Just recently, I read a paper by a biologist called Michael Levin who formulates an idea that's actually very similar to my own idea in that respect. It is that if you go up in evolution to more and more sophisticated organisms, and then you look at the horizon in space and time, the things that they can either remember, or imagine, or perceived to be things that have happened or could be happening.

As you go from a bacterium to a multicellular organism, to a simple animal, to a human, that horizon expands. So we become more and more aware of things that are happening or maybe happening not in our immediate spatial-temporal neighborhood, we become more and more aware of things that may be far removed in time and space. That could be a simple interpretation of this law of growth of complexity.

The field of complexity that means that the things you can be conscious of, tends to increase with evolution and there's a good reason why that should be. That is that evolution is based on natural selection. Natural selection means fitness, means the ability to deal with all kind of problems and challenges. The wider your horizon is, the better you can see any possible challenges coming, the better you will be able to deal with the difficult ones and to protect yourself from them, and the better you also be ready to deal with a positive one to exploit the opportunities.

I think this view of the superorganism is telling us that we should expand our horizon of consciousness from the individual human to the level of humanity as a whole. That will create complexity because a superorganism is more complex than a simple organism, but that will also increase our consciousness. For example, now we are conscious of what's happening in Afghanistan. Hundred years ago, that never have happened.

Now, we should of course be careful not to just say that the wider the spatial range, the more conscious. We can be very much aware of things happening in Afghanistan that are not at all relevant to the future of humanity. There's of course, also the quality of the consciousness, which is something I have also been doing research upon, which is a little bit more difficult to explain in a few words.

I think we can unmistakably say that evolution goes together with this increase of complexity that we go from organism to superorganism and the increasing consciousness in the sense that we become aware of more and more things that are not immediately in our personal neighborhood.

**DSW:** That's great, Francis. That was very eloquent. Let me add my own perspective on that from the standpoint of evolution. Going all the way back, basically the pre-Darwinian notions of evolution such as

Herbert Spencer who you've already mentioned were of course progressive. This is part of the enlightenment movement. We get people like August Comte with his religion of man and Herbert Spencer. These were all more or less secular worldviews that were very value-laden. Were functioning as worldviews, were functioning in the same capacity as religion.

Then came Darwin and Wallace, and this amazingly simple concept, of variation, selection, and replication as the mechanism of evolution. And with the rediscovery of Mendel's work, that led to a period of evolutionary biology and theory which we can look back upon as being amazingly constricted, that all the purpose was drained away from evolutionary theory. Now, it was just the organisms vary and only the immediate environment does the selecting. So, the whole study of evolution went away from Teilhard's vision.

Teilhard's vision was one in which he saw it as a metamorphosis of the Christian religion. Evolution is a metamorphosis of the Christian religion. Then, what we can say now which I think is so exciting and amazing that by going back to basics as I put it, and at seeing Darwinism as any process that combines the ingredients of variation selection and replication, which I think is how you began this conversation.

Then that brings back in cultural evolution and has ample room for a conscious component of cultural evolution. All of a sudden, we're back to where evolution can provide a narrative which is similar in spirit to Spencer and very similar to Teilhard. It can be a science-based worldview that can inform us, along with complexity. So evolution and complexity, the new foundation for our third story.

Although, I'm sure that portions of our conversation were going to be over the heads of a lot of the people listening to this, but the third story is actually something that can be deeply intuitive, commonsensical, and anyone can understand. So, I'm so happy to be a part of that to the best of my ability.

**FH:** I wanted to maybe add an argument for this progressive evolution. It's an argument that I developed in a paper more than 20 years ago which was quite well-cited by way of anecdotes. My collaborator Shima, she discovered that paper when she was making her PhD. That's what brought her to come and work with me. So the paper is called The Growth of Complexity During Evolution and that was controversial idea at the time because you had this approach among evolutionists, among which Stephen Jay Gould, that evolution is just adaptation to local circumstances.

There's an Ice Age. In order to survive, you need to have a thick fur. So, you evolve into something like a mammoth. Then the Ice Age goes and it becomes warmer, so the mammoth needs to evolve into something with thin skin. Then, the climate goes cold again, so the thick skin goes again. It's just random adaptation to whatever vagaries that are in the environment.

So the argument I made is that, yes, it's good to show an adapted to your environment but since the environment anyway is variable, the more adaptability you have at the moment, the better you are. So, there is a tendency to evolve into systems that can adapt to a wide range of things. If they can adapt only to a particular range of things, they are likely to be eliminated by the next change in the environment.

So this increase in the range of things we can deal with, that's something that can be expressed even in the form of a cybernetic law. It's called the Law of Requisite Variety. The cyberneticist actually says, "If you have to deal with various perturbations, the more variety of actions that you can do, the better you will be able to deal with all these problems."

So there is, you might say a selective pressure for living systems to increase the variety of things they can do. In increasing the things they can do, that means on the one hand getting the physiological anatomical features to do these things which means increasing complexity, but it also means a kind of a cognitive increase and increase in the number of things that they can sense and the amount of knowledge they have to know how to deal with these circumstances.

So, the selective pressure is to increase the range of challenges you can deal with and that is a progressive evolution. There is no doubt about that. The example I gave, supposed that you have two systems. One that can deal with situation A, B, C, and the other one can deal with situations A, B, C, D. Well, as long as the situation is A, B, or C, they're both equally good. The moment that the situation changes into D, the first one gets killed off and the second one remains.

So therefore, there is a selective pressure to be able to deal with a wider range of things and that implies all this progressive evolution towards more intelligence, more consciousness, more flexibility, more complexity, et cetera. So, the third story is really that evolution is progressive. We learn to deal with a wider range of situations. We expand our consciousness to a wider range of challenges and opportunities.

**DSW:** That's so great. There's a story that's emerging for human evolution is that it did take place in a period of extreme climatic instability, very prescient for our current times. That evolution of flexibility, Francis, does cut both ways because if it turns out that the environment does become stable and it's always A, B, C, but not D, then plasticity that it includes D is irrelevant and pretty soon does not go to extend in the D direction. So, the degree of plasticity is going to respond to the degree of environmental variability.

That said, I think the story that's emerging is that actually during this period of climatic instability, all mammals, their brain size increased over evolutionary time. So, they all became more plastic. Humans of course became plastic in a new way their capacity for cultural evolution, which by the way is inherently cooperative. So human cooperation and human adaptability I think are joined at the hip.

Francis, I wanted to finish up with one final point which I think that so much integration is taking place between, and needs to take place between evolutionary theory and complexity theory. Those two bodies of thought are only now going together. A major difference I think that I encounter again and again, has to do with whether the emergence of the noosphere, the global superorganism is something that's just going to happen bottom up or whether there needs to be some more deliberative process of selection which is not the standard engineering kind. We both agree upon that.

There are some sense that if we want the earth to function as a whole system, we really have to be very deliberative about what we select with the whole system in mind as opposed to that this will be that's not needed so much and this is something that will just happen. I'd love to know your thoughts upon that as perhaps our final topic that we consider.

**FH:** Well, even though I mentioned the invisible hand, I'm not really a believer in laissez-faire economics, and the same applies to governing society as a whole. You need to leave space for all these bottom up processes of self-organization but in many cases, it's necessary or at least it's more efficient to also have some top-down processes.

The point is that these top-down processes you have as you said before, you have to be able to experiment. You should not assume like we know best. That's the way we are going to organize society. That's the way it's going to work. No, you need from the top-down also experiment and say, let's try this policy, for example. Now we have for example, the problem of what do you do with people who don't want to get vaccinated against Covid. At the global level, it's best that everybody should get vaccinated, but there are some quite big groups that for various reasons don't want to get vaccinated. Then, you can experiment.

You could for example say, "We will make it obligatory that you need to have a vaccination to work in certain places or we will create privileges for people who are vaccinated that they can do more things, or we can try to make it easier for people to get vaccinated." All of these are strategies that from a top-down point of view may work but you don't really know whether they will work until you apply them.

I think that is also one of the approaches we need to think about this. This top-down thing is useful. It's important that you need to think and to experiment with it because there is no obvious way to do it.

**DSW:** That's great, Francis. That embodies Prosocial's philosophy which we call bottom up meets enlightened top-down. So, that's the way that we put it. Okay. Well, I think we're done. I think this has been a great conversation. I've learned a lot from it, Francis. It's been really one of the nice things about this project with Ben is that my opportunity to get to know you and your work much better.

**FH:** I had already known about your work, but not as much as details as I should have. So, I'm also happy that I get an opportunity to get to know your work better. I hope that within the project, we may be able to come closer. I think we were kind of coming closer. The whole issue is the selection at the global level versus the selection at the local level, how much of each. I think we can come to some solution.