Francis Heylighen with David Sloan Wilson, Part Two

David Sloan Wilson: 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?

Francis Heylighen: 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, Max Wildiers, 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 connected 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 something 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 a 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 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.