NETWORK LAW REVIEW

ARTICLE

Some Background to Complexity Economics

W. Brian Arthur

***

The following article is based on a transcript of the keynote talk given at the DCI conference by W. Brian Arthur on April 13, 2023. It has been approved by the author. For further reading, see "Foundations of Complexity Economics," W. Brian Arthur, 2021, Nature Reviews | Physics, 3, 136 – 145.

Brian Arthur: I'm delighted to be here, even if it is pretty early in the morning here in California. What I want to tell you is my version of what I think complexity economics is about. I will not go into technical details, but I want to give you an overview. I'm sure most of you have heard of complexity economics, some of you may even be practitioners. Here is my take on it.

I will start with the standard neoclassical economics that I was brought up on. It is very much based on mathematics, and to make the mathematics work we are forced to make many simplifying assumptions. In nearly all models that we look at with standard neoclassical economics we are using equation-based mathematics, and we tend to assume that all agents are identical to keep things simple. So we assume that the agents are all the same, and further we assume that they have perfect knowledge of other agents, that all decision makers in the economy are perfectly rational, and that together, all agents arrive at optimal behavior that is consistent with, or in equilibrium with, the overall outcome caused by that behavior. This is an ingenious shortcut to doing theory. I'm dazzled by the mathematics, it gives you a view of the economy that is extraordinarily elegant. I'm trained mostly as a mathematician, so I find this quite gorgeous. But – and this is a huge but – it is highly restrictive with its assumptions, and very often unrealistic. I'm not the only one who says that. For about 100 or more years, economists have been complaining: 'It's beautiful, but hang on, just how realistic is this?'

Complexity economics relaxes these assumptions. It comes out of a thought process: What would it be like if we had an economics where in the models we build,

agents could realistically differ? What if they didn't have perfect information about other agents? What if they just didn't know much about other agents? This would leave the agents trying to make sense of the situation they are in. You would see them exploring, they would be reacting to what they see, they would constantly change their actions and strategies, and update these as well as they could. That may not lead to an equilibrium; equilibrium is not assumed. If there is a natural equilibrium in that situation, it might emerge, or it might not.

Once you start to assume more realistic assumptions, the problem is that models start to get more complicated. When they get more complicated, we have to resort to keeping tabs computationally. If I have a million and one or 1001 agents, and they differ, I cannot keep all those in my head. I need to resort to computation.

I want to say a few words on where this framework came from. It did not arise out of nothing. There was a meeting in the very incipient Santa Fe Institute in New Mexico in 1987, convened by Kenneth Arrow and Philip Anderson. Anderson is a top Nobel Prize winning physicist and Arrow, as I'm sure you know, is a top economist. Arrow brought 10 theoretical economists to the meeting and Anderson brought 10 physicists and other scientists. Anderson's group included physicist Doyne Farmer, biologist Stuart Kauffman, and computer scientist John Holland. The theoretical economists included Tom Sargent, Larry [Lawrence] Summers, Buz Brock, and other people you'd recognize, as well as myself.

The meeting was spectacularly successful and went on for 10 days. We were all dazzled by what the physicists were doing, and they were a bit dazzled by us. They thought most of the neoclassical stuff was a little bit old fashioned, but it gave them an insight into economics. The Santa Fe Institute subsequently decided to constitute a program of research – the Economy as an Evolving Complex System. A year later, I was brought back to head up that program. This was the Santa Fe Institute's first program. We didn't really have staff; we barely had a building – the Santa Fe Institute at that time was a startup.

Once the program started, in August 1988, we had economists and physicists. We found ourselves sitting around a table in the kitchen of an old-fashioned convent we had taken over. We didn't quite know what to do. There was a kitchen where we could sit around and debate, 'We have money from Citibank, we have the backing of top economists and physicists – what on earth are we going to do?' This debate went on for three or four weeks and as leader I wasn't at all sure of what direction to take.

In the end I called Kenneth Arrow at Stanford. This was before emails. 'Ken', I said, 'We are not quite sure what to do or how daring we can be here. What do you think?' Arrow called Anderson at Princeton, and Phil Anderson called John Reed, who was chairman of Citibank at the time and was putting up the funding. The word

came back from Reed via Anderson via Arrow to me, 'Do anything you like, providing it is at the foundations of economics and is not conventional.' I was staggered. The thought I had immediately was that, suppose in 1520, Martin Luther had indirectly got in touch with the Pope saying that 'We are redoing theology, what would the Vatican recommend?' and the word comes back from the Pope, 'Do anything you like, providing it is not conventional and is at the foundations of theology.'

We decided to drop the whole notion of equilibrium. We were not against equilibrium, but that wasn't going to be an assumption. We decided that if we built models, that the agents in these models could differ. Then, of course, we realized that we were in hot water immediately, because, if agents could differ, we should also assume that other agents didn't know exactly how they differed and how they thought. That was subject to fundamental uncertainty — you simply don't know. I'm standing here talking in Silicon Valley and if something new is getting launched here, you don't know what your rivals are thinking – not in any detail anyway.

So we faced a problem. The moment you introduce fundamental uncertainly, your problems become ill defined. There is no amount of logic that can solve things if certain parts of your problem are fundamentally uncertain. You don't know what the problem is. Once the economic problems are ill defined, solutions are ill defined too, and you grind to a stop.

We had John Holland with us, and I think, to this day, that was a miracle. John had been spending decades teaching computers how to get smart when they didn't know what situation they were in. We borrowed a lot of his thinking and a lot of his methods, and we realized that in the real economy, agents may not know totally what they are doing or what problem they are in. Think in terms of three years ago when we didn't know what COVID was, we didn't know how severe COVID would be, we didn't know when vaccines were coming. We didn't know. Yet we all went forward together, albeit a little bit scared. And so in this way we could model agents learning along the way, mutually, as the situation unfolded. This gave us an overall approach and we developed it by continually solving standard problems using this new outlook. There were other groups working in parallel – Rob Axtell, Josh Epstein, Alan Kirman, and others – but the main effort, I would say, was at the Santa Fe Institute.

At the end of the decade, I was asked to do a paper in the journal *Science*. The editor called me from London and asked, 'What do you call this new approach?' I said, 'I don't call it anything.' He said, 'No, no, you need to give me an answer. What do you call your new approach?' This went back and forth and eventually I lost. I said, 'All right, call it complexity economics.' It was one of these things dropped from heaven. I was standing with a landline in my apartment in Palo Alto and that label locked in.

If we have a different approach to solve problems, in any science, a really good question to ask is: Can it solve any problems that the previous approach cannot, or can it give more satisfaction in explanation than the previous approach? This really concerned me. So, we selected a puzzle in standard economics, it is called the question of asset pricing. Think of it as a question of how stock market prices in a simple stock market arise. If you have some earnings that go up and down stochastically, how would prices track the earnings? That problem had been solved by Robert Lucas in 1979 using standard economics, standard assumptions, and good mathematics. The Lucas solution, I thought, was dazzling. It was brilliant, gorgeous, and elegant work. But he did assume standard assumptions, and in fact, he assumed all investors were identical. So, there's a 'but' to this.

The Lucas solution very much looked like, and tracked, how real markets worked. I thought that was a huge step forward, and beautifully elegant mathematics. But there are several phenomena you see in real stock market pricing that the Lucas model couldn't show and didn't show: price bubbles and sudden crashes; technical trading (meaning agents pay attention to past prices and to trade volumes); periods of high volatility followed by periods of quiescence; a market psychology where investors have different ideas of the market. Lucas's solution showed none of these phenomena. It also showed, embarrassingly, that in this market trades are zero. That sounds shocking, but if everybody is identical, then they all either want to buy or they all want to sell. They can't trade because there is no one on the other side, and the price adjusts so that everybody's indifferent between buying and selling.

In 1988, John Holland and I put together a team and we set up Lucas's model to work on a computer. We slid out Lucas's identical investors and replaced them with ones that in fact could differ. We made each of these investors its own little computer program, a tiny little computer program that could make decisions. Each investor, call them artificially intelligent investors, could respond to its own collection of conclusions – If the market is doing *this*, then I forecast *that*. If the markets doing something else, then I have a different forecast. Rules could be thrown out if they didn't work well, and new ones could be developed as necessary. In other words, we set up what is now called the Santa Fe Artificial Stock Market, with investors who were small primitive AIs, and let them make bids and offers for stock.

We set all this up around the Lucas model. We had a bet with our good friend Tom Sargent, an excellent theoretician. Tom said, 'You know, you're not going to get anything out of this weird and wondrous model with different investors. It will just be attracted to the Lucas solution. You're not going to see anything new.' When we got the price series, sure enough, it looked very similar to the standard neoclassical solution of Lucas. It seemed Tom Sargent was right and I was really disappointed.

Then we looked more closely. We worked out the standard Lucas solution and plotted our model's price series against that. What we saw with this closer look was quite amazing. We saw bubbles and crashes, we saw that our model showed the emergence of technical trading, we saw periods of high and low volatility, a market psychology emerged, meaning different opinions, trading volume was not zero, in fact it was quite significant. All of these phenomena you might see in real markets in New York, London, or Frankfurt, and they surfaced in our model as emergent phenomena. Our new approach had revealed real phenomena that the standard methods couldn't see.

There have been many studies since – some in Santa Fe, many elsewhere – and one overall theme that we see is that the economy that results is not so much 'a machine that links this with that, that affects this, and equilibrium is reached.' It is much more like an ecology. What you are really building is an ecology of different actions, different beliefs and forecasts, and maybe different strategies. As this ecology evolves and changes, different beliefs and strategies get emphasized. In general, there is no equilibrium. You might see perpetual novelty, but within the system, agents learn and adapt. Temporary phenomena, like very high volatility or low volatility, may emerge, and maybe even new behaviors emerge and are discovered. Under complexity economics, the economy becomes something not given, not preexisting, but constantly forming from a developing se t of actions, strategies, and beliefs. It is not mechanistic, static, timeless, and perfect, but rather, I would say, it is organic, always creating itself, and alive, brimming with messy vitality.

Why did this approach arise now – why not in 1939 say, at the time of where Paul Samuelson was giving deep thoughts as to foundations of economics? This approach of making things more realistic and of trying to see how things play out more realistic, is not new. A lot of economists 100 years ago must have thought of that. What is new is the coming of computation that allows us to look at these more realistic, complicated approaches.

I'd like to finish were with a comment. New tools in economics always bring new theory. When geometry came along, we got a lot of black-board type theory, but it gave us new insights. When algebra and calculus came along in the 1870's, in due course by Samuelson's time that brought us neoclassical theory. I think we progressed a lot in getting neoclassical theory. Now, computation is bringing us an economics that can handle heterogeneous agents, and handle fundamental uncertainty, because we can model how people actually behave when they don't know – and it brings us non-equilibrium. It is not surprising that this approach arose in the 1980s and early 90s, because that is when we all got desktop computers.

Thank you very much.

——-

Q&A

David Teece: Brian, thank you for a brilliant talk. I think the members of the dynamic competition initiative are very sympathetic to complexity economics, and systems theory more generally. But give us two or three principles that we should keep in mind, as we think about what I call workable systems theory or administrability, if we try to apply these principles to competition policy and innovation. Are there one or two things that we should always bear in mind as we think about antitrust enforcement?

Brian Arthur: I realized yesterday that I was talking to a group of people, particularly David here, who have developed a dynamic framework for looking at competition. I would say this, generally about the approach: Probably what is missing here as yet is - and I'm sure people are working on this, like Thibault [Thibault Schrepel] and Bo [Bowman Heiden] - what the relation between these two different approaches is. I don't see that they are very different.

One comment I would make, if you are looking for principles, is that dynamics, or in our case complexity economics, are not assumed. We allow agents to react to the situation they are in. That reacting brings about change, and change brings in dynamics. One of the assumptions we dropped in complexity economics was the assumption that there were always diminishing returns on the margin. In fact, as you all well know, when it comes to tech companies or technology, very often are they increasing returns on the margin. The more of my friends join PayPal, the more likely am I to join PayPal so that I can transact with them. As PayPal gets larger, it gets more advantage.

I think that what emerged in the 1980's and 1990's, was an awareness that you could do economics in markets that showed increasing returns and therefore lead to large concentrations of market share — lead to monopolistic situations. You could do that, if you model this as small events probabilistically leading you into one of many outcome. That goes back to what Alfred Marshall said in 1891, that if you have got *N* firms competing under diminishing costs, then the larger they get the lower their costs go. The market, he said, will go to one firm, but you can't in advance tell which. I think that one strong connection, or principle, is that you don't need to get scared of multiple equilibria. Sometimes one equilibrium is reached, sometimes another one, and that leads me into thinking that in all these cases, not just the dynamics emerge, but the outcome emerges as well.

The use - and I make it a third principle here - is heavily dependent upon computation. Strictly speaking, computation isn't necessary. All the early work I did in this field was done by using nonlinear stochastic process theory. You can do

mathematics if you are willing, but in general, to keep track, to theorize and to think deeply about things, we are forced to look at situations computationally. Only then can we figure out how something emerged.

My overall answer is that I wish I knew much more about dynamic competition, and I think there is a strong relation between the framework I'm proposing and the framework that you are dealing with here. All I can do is cheer you on.

Gordon Phillips: I'm a big fan of using computers in economic research, but there is one issue that I always struggle with. At some point, you have got to reduce some of the complexity. We can't have 5,000 different consumers of different types. So, how do you put bounds on this if there is no equilibrium force that occasionally winnows out some irrational consumers or irrational actors?

Brian Arthur: Great question! Computationally, it is not necessary to say that there is an upper limit like 5 000 agents. To my staggering astonishment, my good friend Robert Axtell produced a model recently – I think it was in the housing market in the UK – where he had several million different agents in his computerized model. You try to make sure the types of agents you get do reflect reality. You can do that. Or maybe you are just using a thought experiment.

There is a subtle point here about rationality. Once you don't know what other people are doing, the problem becomes ill defined. If the problem is not well specified, economically, if we just don't know, Keynes talked about that sometimes we can put a probability on the situation. He said that we don't know whether there will be a war in 1937, and we don't know what the price of copper will be. You could put probabilities on there and define a model. In what we are doing here, we don't resort to that. We assume we just don't know. Agents are looking, as if they are exploring the dark, feeling their way, trying to figure out what tunnel they are in underground and where it is leading. If the agents don't know and they cannot optimize, there is no such thing as a rational agent. If there is no rational problem, there cannot be a rational solution. We make no assumptions about rationality.

The results I gave for the stock market are not due to irrationality, because they are the result of exploratory agents finding ways to learn and cope in a situation of other exploratory agents finding ways to learn and cope in that situation. Rationality might emerge, an equilibrium might emerge, but it is a very different way of thinking. It is a wonderful question, thank you.

Johannes Bauer: Thank you very much. In the session right before this one, we discussed a European new legal initiative, the Digital Markets Act, that aims at improving the governance of digital ecosystems. I'm not assuming that you know the details of this law, but in general, from your research on ecosystems and complex economic systems, what is the role of governance? Can governance improve the

outcomes of those systems? How can we model this, and analytically try to grasp how governance influences the dynamics in such complex systems?

Brian Arthur: I think there has been a lot of work here, and what people are talking about is setting up policy labs. What you are looking at is not so much to make a definitive model, an equilibrium model, or a rationally based model. Rather, you are trying to set up a model that captures the situation you are looking at. You can regard that as a lab experiment, to see what the outcome is. You can change certain parameters in the situation as you like, and then rerun the model to see what the outcomes are.

It is not, I think, a very deep approach. John Holland used to talk about flight simulators for the mind. If you can simulate an aircraft, you can put it into different situations and see how it responds, and I think that is what a lot of people are doing for policy. I don't think this is particularly new, but it is a different flavor from neoclassical methods.

One thing I would point out is the analogy of an ecology. I would say that the standard neoclassical approach is very much an engineering one, where you regard yourself as being at the controls of a very large power station, and you are turning this dial for policy. Once you start to think of the outcome as being an ecology of competing, or sometimes cooperating, behaviors, then you can think of yourself not so much as tweaking the system with large dials, but more as a park ranger – I'm introducing this type of species, I'm fencing here, and I'm looking after things there. For me, that is a very different outlook. In real ecologies, an awful lot of trees were wiped out in the 1930s and onwards in the US, and no doubt in Europe as well. Because it became legal to shoot wolves, wolves disappeared. That meant that elk and other tree grazing animals became plentiful. That meant that small fir trees disappeared, and that changed the whole ecology. It is not radically new thinking, but it is a different way of approaching things.

Thibault Schrepel: Brian, I have a 'complex' question and I'm going to ask you to give me a short answer, which I know is unfair.

With agent-based modeling, you have a way at the very micro level to assess how each agent may address tradeoffs in a very different manner. On the macro level, you see that agents learn from which of their hypotheses work and create a macro trend. This was the point you made in your El Farol paper [Arthur, W. B. (1994). Inductive Reasoning and Bounded Rationality. *The American Economic Review*, 84(2), 406–411.

If we were to develop those agent-based models to test the impact of competition law, or any regulation, what will be the key element to keep in mind, so that the result of the simulation is not totally bonkers?

Brian Arthur: You are right, that is a complex question, and I can't find a simple answer. My instinct would be to do two things, and they are somewhat at odds. I would try to model the situation fairly simply, then run it and rerun it under slightly different circumstances and slightly different assumptions until, looking at the outcome, I felt that I had a very good intuition for how that worked. Then, I would try to look at reality and say what are its characteristics.

In the United States, there is a debate in the market for Uber and Lyft, discussing whether there are increasing returns, and if there is a lock-in. One thing you can do, is to make assumptions and try to get a feel for that whole market. It seems to me that the increasing returns early on increase, then flatten out. Doubling the number of Ubers around isn't going to help that much, generally. I would explore the situation until I fully understood how it worked and had a good intuitive feel from looking at, say, agent-based models. Then, I would look at reality to see where it does fit in the parameter space of what I have been looking at and get a good feel for that. I hope that is, more or less, where you are heading.

Thibault Schrepel: Brian, thank you so very much. Please join me in applauding our keynote speaker.

W. Brian Arthur
External Faculty Member at the Santa Fe Institute, and Visiting Researcher in the Intelligent Systems Lab at PARC (formerly Xerox Parc).

***

Citation: W. Brian Arthur, "*Some Background to Complexity Economics*", Network Law Review, Summer 2023.