The economic entomologist:
an interview with Alan Kirman

ALAN P. KIRMAN (London, 1939) is professor emeritus at l'Université d'Aix-Marseille III and researcher at GREQAM (Groupe de Recherche en Économie Quantitative d'Aix-Marseille). He has published over a hundred academic articles and edited and authored many books including noted monographs on general equilibrium analysis (Hildenbrand and Kirman 1976; 1988) and most recently *Complex economics: individual and collective rationality* (Kirman 2011, reviewed in this issue of *EJPE*).

Professor Kirman's work touches on many issues of interest to economic methodologists and so we were delighted to have the opportunity to interview him when he visited the Erasmus Institute for Philosophy and Economics (EIPE) in late November 2010 to present a paper on the state of macroeconomics.

In this interview Professor Kirman discusses his understanding of the relationship between individual behaviour and aggregate patterns, why it is essential to consider the interactions between agents, and what the study of ant's behaviour can teach us about collective human actions. He explains the core concepts of his 'interactionist' approach, including microfoundations, rationality and emergence, and reflects on the potential of agent-based modelling, the limitations of game theory, the possibility of aggregate-level analysis, and the relevance of behavioural studies. The interview also ranges more widely, discussing the different goals of economics (for instance, explaining, predicting, and controlling), the role of mathematics in modern economics, and the state of macroeconomics.

EJPE: *What brought you to economics in the first place and how would you describe your research in the early years of your career?*

ALAN KIRMAN: Well, my story is a bit weird because I started out after my first degree at Oxford as a school geography teacher. But I found myself asking: "Do I want to do this for the rest of my life?" I went to some evening classes in economics organized by the Workers Educational Association, and I thought: "Oh, that is really interesting. Maybe I should try and do something with this."

I first did a one year diploma (part of it on international economics) at the Johns Hopkins School of Advanced International Studies in Bologna (Italy). One of the people who were teaching there, Ira Scott, gave me a recommendation to do a PhD in Minnesota, so I took off to Minnesota the year after. There, it was very cold and extremely mathematical. But I had done no math before, so my advisor Hugo Sonnenschein told me I had to do a degree in math as well.

By spring, I thought that this was not what I came there to do. For me, economics was about unemployment and inflation and so on, and yet here I was struggling with fixed point theorems and all that stuff. I said to Hugo: "Look, why couldn't I do these other things?" He said: "No, no, that is macroeonomics, and macroeonomics is about wisdom. Microeconomics is about analysis, and young people should do micro".

Anyway, I got a fellowship to go to Princeton, where I thought I was going to do more *real* economics. Once there I looked around at the people who were teaching and by far the most interesting and inspiring teacher was Harold W. Kuhn, who was—unfortunately for me—a professor of mathematics and also a professor of economics! So I did my thesis with him, applying non-cooperative game theory to international trade.

Afterwards I moved into general equilibrium. I worked on a lot of other things too in my early years, because, you know, general equilibrium is not very inspiring. I mean, it is a great intellectual game, but it is so mathematical. So I worked on fairness, social choice, a bit on international trade—lots of different things, nothing very deep, and that is how I started out.

Can you give us a short 'insider' story of the growing dissatisfaction with the general equilibrium project—the role of the Sonnenschein-Mantel-Debreu results for instance?

The results of Sonnenschein (1972), Mantel (1974), and Debreu (1974) had a major impact. In his book *Market demand*, Werner Hildenbrand (1994) said that he used to work under the assumption that the microeconomic foundation would give us predictions at the aggregate level, but these results overturned that.

We learned that more or less anything goes in terms of what we can observe at the aggregate level. For example, how can you possibly do comparative statics in macroeconomics if you have several equilibria —which equilibrium do you start with and which one do you go to? That is one problem. Also, if you have no idea whether economies starting out of equilibrium will actually go to equilibrium, why would you be so interested in equilibrium? Michio Morishima (1984) said that such a concept would be of purely intellectual interest.

So these are the sorts of things that made people uneasy about the very restrictive conditions we typically impose on individuals. When you learn that even with those conditions you are not going to get much structure in the aggregate, you start to say to yourself that maybe this project has some sort of fundamental weakness. A lot of microeconomists said that this was not very good, but macroeconomists did not take that message on board at all. They simply said that we will just have to simplify things until we get to a situation where we do have uniqueness and stability. And then of course we arrive at the famous representative individual.

Hildenbrand took a radical stand and said that we should forget about all this individual optimization and just look at distributions of how people behave and see if we can extract rules about the economy. That went over like a lead balloon! He gave talks at Berkeley and everywhere and people were just extremely hostile. So his approach never took off and yet I think it is a very positive way of doing things.

Your recent book is called Complex economics (Kirman 2011). Many economists in the last twenty years or so have endorsed similar labels. What does 'complexity' amount to? How did you shift from general equilibrium to this other project?

The *Journal of Mathematical Economics* was first published in 1974. Hans Föllmer—a mathematician in Bonn at the time, later in Zurich, and now in Berlin—had a paper in the first issue of this journal which was called "Random economies with many interacting agents" (Föllmer 1974). He showed there that if you have lots of people who have their

regular preferences and so forth, but those preferences are influenced by their neighbours—like particles in the Ising model from physics—that could destroy the underlying notion of a unique equilibrium or, put alternatively, that you could not say much about the aggregate once you had this interaction.

I started to think about interaction models and talked to Hans about that but I was not imaginative enough. Then I met a mathematician from Warwick called David Rand and we had a long discussion about whether we could think of demand differently, with each individual's demand being influenced by the demands of his "neighbours".

I think this was when I started to think of systems where you have these really quite primitive individuals interacting. That is what came to underlie my view of complexity: lots of rather simple individuals who by their interactions generate phenomena at the aggregate level that do not coincide with what you see at the lower level. In economics typically, we make a short-cut: we just assume that what is going on up there looks like what is going on down here.

The term 'microfoundations' is highly popular in economics. Would you say that your own approach is critical of the quest for microfoundations or that it is only critical of the specific 'microfoundations' in vogue in economics?

It depends on what precisely you mean by microfoundations; do you mean by that that I want to reduce the economy in some sense to looking at some sort of 'typical or representative individual', so that once I understand how that one individual works, then I understand how the economy works? Well, that I am critical of (see Kirman 1992). I am also critical of the particular microfoundations that we use. Why should we focus so closely on these axioms of rationality *which we have imposed*? Those are the two aspects I would be critical of.

This does not mean that you should not be interested in what individuals are doing—after all, that is what agent-based modelling is about. You can also be interested in how individual behaviour affects the working of the economy but you should not have a simple, add-up, clockwork idea of this transition.

Somewhere in my new book (Kirman 2011, 19) there is a quote from two neurologists, Churchland and Sejnowski (1992, 4). They say that aggregation in neuroscience is much more complicated than simply summing up the properties of neurons. This should not stop you from

being interested in what is happening at the neuronal level but, if you want to understand people's attitudes and what they think, you are not going to stick to examining neurons, right? You will be interested to know what happens to the network of neurons—how they combine and which ones get fired when you are in certain moods and so forth.

You reject the standard rationality assumptions in economics. What are your reasons for doing so? What are the substitutes?

Those axioms—and you can find a whole series of people from Pareto onwards who make the same argument—come from economists' introspection and what they think is necessary for their work, not from observation of what people are doing.

Some of these axioms seem natural, at least at first sight. For example, transitivity seems a natural idea—if you prefer A to B and B to C, you also prefer A to C. But if you look carefully at how economists define the things over which you are making choices, you could never observe whether or not an individual is making transitive choices.

My main problem is that none of these axioms is taken by observing lots and lots of people. In other disciplines, that is what you do. You look and then you try to develop a model which might explain what you observe. In economics, we started out by doing the formalization and building models which were internally consistent but often far removed from reality. To construct models which we could analyse formally, we needed to make some formal assumptions. As I said, many scholars starting with Pareto basically made the same remark: these assumptions are somehow not natural, they are not about what people do, they are more about what we need in order to pursue our analysis. So that is my real objection.

What do you replace that with? Do you just say that people just make arbitrary random choices? Well of course not. The argument I would make would be that, in some sense, people see directions in which they think their welfare improves, and they try to move in those directions. A simple way to model this is to give simple rules to agents that you find plausible and then look at how that works. In such a model, people are not *irrational*, but rationality must have a much more open definition.

You display some sympathy with the project in behavioural economics to supply psychologically-refined assumptions regarding economic agents. At the same time, you assert that we have much to learn from studies of ants, bees, and other social animals which can be modelled as acting based on simple behavioural rules. Are these two lines of inquiry—refining the psychology of our modelled agents and looking for simple behavioural rules—not in tension?

The distinction is really the following. Say that I observe people behaving in certain ways and that together they are generating some aggregate pattern. I say to myself: "Can I think of this phenomenon in terms of the rules people are following without worrying about their intentions for the moment?" That is like the ants phenomena, in the sense that the ants are interacting in very simple ways.

Now you might say that you are interested in why they act like that. Why are they following these rules? One can say that evolution has led them to select the rules that they follow. This is what is often said about ants. Yet Deborah Gordon, a famous entomologist has collected a mass of evidence (see Gordon 2010) to show that quite often ants are individually inefficient and fail to do what they are trying to do. Although they achieve a lot, they do not seem to be behaving optimally in any standard sense. I think economists have bought in too easily to the unsophisticated evolutionary arguments. Her advice to people who argue for optimal behaviour is "spend time watching ants". The same advice could be given to economists, "watch economic agents!" Thus, when we are looking at human beings we probably want to know much more than that they seem to follow, in general, simple rules and we want to look at the psychological side of things. That is why I have been interested in neuroeconomics and published a couple of articles on that (e.g., Oullier et al. 2008; Kirman and Teschl 2010).

So I do not think there is a contradiction. When we are interested in humans we are not only interested in knowing which rules they follow but also why they follow these particular rules and why they often seem to behave non-rationally. This would be my distinction.

You have made many contributions to the difficult question of how to connect macro-patterns to the behaviour of agents. What are the main results regarding this question? For instance, can you tell us about your fascinating work on fish markets?

There are two important results on the relation between the aggregate patterns and individual behaviour. Firstly, even if individuals are not 'well-behaved' in the standard sense, the aggregate can actually have rather nicely orderly behaviour. Second, it is also possible that when we have a lot of well-behaved individuals, the aggregate seems to be much messier than the behaviour of the individuals.

In considering the first possibility, that is where I go to look at real markets—like fish markets (Härdle and Kirman 1995; Kirman 2011, chapter 3)—and try to understand what is happening. The individuals in markets are in very complicated situations; they know a lot of information about all the people around them, and various things are happening which are not usually incorporated in our models.

You could try and model this as a huge and complicated game. But, in fact, these people settle down to doing things that are really rather simple. For example, a lot of them become extremely loyal; they always buy from one seller. Once you have established that fact, then you can begin to see a sort of pattern emerging in the market. Some people are shopping around, but overall you get a rather nice negative relationship (which to be precise I should not call 'demand') between the amount of fish on the market and the realized prices. The simple patterns of individual behaviour actually generate at the aggregate level something which looks quite like what you might have thought of as a classical aggregate excess demand function, but it is certainly not derived from every individual having his own regular excess demand function. (Incidentally, there has historically been a lot of discussion in economics as to whether what we observe when people make purchases corresponds to the formal definition of demand! But that is a discussion for another day.)

Considering the second possibility, if you look at financial markets (Kirman 2011, chapter 11), people who seem to be following rules which are in some sense rational can sometimes generate huge movements in the market, not because these people suddenly become irrational but because they are taking into account what other people are doing and so they start to follow the trend. Information cascades and these sorts of phenomena arise not because people are just weird. They happen because people interact with each other. This interaction can lead to mass movements which you could not have anticipated if you just looked at individuals in isolation.

So you get two things: 'bad', 'irrational', or 'odd' behaviour which somehow gets smoother, and the other way around.

The term 'collective rationality' is in the subtitle of your recent book. What do you mean by it? How does it relate to our usual understanding of individual rationality?

Well, I am not happy with 'rationality'. One of the problems that we find is that people have now somehow absorbed the economist's notion of rationality, so that when people say 'rational', they immediately have in mind something like what economists define as rationality. In fact, rationality can be thought of in many different ways.

Rationality for me would mean something more like coherent or interpretable behaviour; behaviour that is not just random. So 'collective rationality' would mean that in some sense this group or society moves in a way that you can observe and anticipate and seems to be purposeful—although I do not want to insist too strongly on 'purposeful' because it is not clear that the aggregate has purposes. In this sense, 'collective rationality' could well be applied to a physical system, where there is clearly no intention involved. Take a system made of physical particles (see Kirman 2011, chapter 6; and Vinković and Kirman 2006). The system's basic tendency is to minimize its total energy. One might want to say that the system *tries* to reduce its energy, but it is not intentional. The system does not have an intention but you can still observe it minimizing its energy and that is something that is well defined. In the end, what I am after is perhaps more a sort of collective coherence rather than rationality.

Some might conclude from a discovery of 'collective rationality' that it is acceptable to simply use techniques that concern themselves with analysing the aggregate level only. Since connecting the behaviour of individuals to macro-patterns is so difficult, such an approach is certainly analytically appealing. What would you say of such an alternative?

Actually, this alternative is in the spirit of the old macroeconomics where we used to have relationships between aggregate variables, and then you have things like Goodwin's business cycles (Goodwin 1951).

It is not an illegitimate activity to think in terms of aggregates. You do not necessarily have to be interested in explaining aggregate relationships in terms of individuals. Central bankers actually often look

at rather simple aggregate relationships without worrying about what it was that motivated people; and they do not even try to derive the aggregate relationships from underlying models. For many purposes (particularly policy purposes), focusing on simple aggregate relationships may even be better than worrying much about all the mechanics of the economy. You may also be interested in the mechanics, but for certain purposes it may be perfectly legitimate to want relationships between aggregate variables. To use a familiar metaphor, you do not have to understand the mechanics of a watch to be able to understand the regular movements of its hands.

You often use the term 'emergence' in your own work. What do you mean by it? Among the many interpretations of emergence in philosophy (see O'Connor and Wong 2009), at least one, the irreducible-pattern interpretation, seems to imply that one would not be able to analyse some aggregate-level, emergent properties in terms of the interactions of units at a lower level (e.g., individuals). But your strategy seems to be exactly that. Why should we expect the study of individual-level interactions to be a fruitful way to analyse aggregate-level properties if the latter are deemed emergent?

I am not a philosopher, so I do not know much about these things, but if you look back, people who were at the interface—J. S. Mill, and people after him—were interested in exactly this distinction between what is happening at the different levels. My primitive, non-philosophical, feeling is that it is not a distinction between looking at the individuals *and then* looking at the interactions between them. What generates the difference at the aggregate level is that individuals are interacting with each other. So I cannot take that individual, examine him (the way he behaves and his decisions), and conclude from that what the crowd will do. I cannot because I am eliminating the essential part which is the interaction. I would say it is the individual characteristics *plus* their interaction which generate the activity up here, which has different characteristics from the specific individuals. It is the fact that one cannot derive the aggregate property from adding up the behaviour of individuals that makes aggregate phenomena 'irreducible', I think.

One main argument—if not the main argument—in your recent book is that "direct interaction between agents plays a crucial role in determining aggregate outcomes" (Kirman 2011, 35). What is so

"crucial" about it? Is it that direct interaction is more "crucial" than other elements—e.g., the behavioural rules of the agents themselves?

Well, when you come to economics, at the start you are told it is about the distribution of scarce resources amongst competing needs, or whatever—they give you a definition. And you say, well how is this achieved? Well, we are told, this is achieved by people trading with each other and collectively that leads to outcomes that have certain properties. But then you say to yourself: "trading with each other, how is that organized?" And they say: "Well, it is through a market. There are some prices which are given and then everybody uses those prices". Yes, but who trades with whom?

In the standard model, the part where people meet each other, trade, and so forth—in which things happen—is just missing. As soon as you start to think about it, you realize that, if people are interacting with each other in markets, what one person is doing will influence others. For instance, when I meet someone and he tells me that he is buying an asset, that would probably influence what I think of its prospects.

All this interaction seems to me important, and yet that is something that we just push under the rug in the standard set-up. There are very few markets where the actual mechanism of dealing—the actual influence of one person on another—is not important. If you want to understand economic activity you cannot lay aside the fact that it has to happen between trading partners.

You seem to be saying that it is crucial for us to look at interactions because they are usually missing from our usual story of how economies work. There would be another interpretation of "crucial" which would be something like the following. One starts with the proposition that aggregate outcomes are determined by the properties of individuals—i.e., the rules they follow—and by the structure of their interactions. The structure of interactions would be "crucial" in the sense that the aggregate outcome will be more sensitive to changes in the structure of interactions than to changes in behavioural rules; changing the interactions a bit modifies the aggregate outcome, while this outcome stays the same for a broad spectrum of behavioural rules. Would you also subscribe to this interpretation of "crucial"?

I think crucial has both these meanings. What we know is that, if you change the rules of the market a little bit, you can actually make a big

change in the result. Let me give you a very simple example. Take a market for a financial asset which is organized as an order book. You put in your orders and orders are cleared as they meet their counterpart, i.e., as the price asked is lower than some bid.

One rule governs for how long those orders stay on the book. Do they stay on there just for a day and then get wiped out? Do you leave them on there for three weeks? Do you leave them on there until the people who placed the orders take them off? Here is a simple organizational rule which does not have anything to do with the intentions of the people who are putting these orders on there, but which changes the actual prices a lot. Changing in apparently minor ways the rules of the organization can actually change what happens in quite a major way. So what you allow people to do with each other, rather than the particular behavioural rules that they are following can actually turn out to be very important.

The typical way of modelling interactions between agents in economics would be to use game theory. But you express some dissatisfaction with game theory. Is not game theory more flexible than you depict it? Do your criticisms apply as well to evolutionary game theory for instance?

If you read Binmore's *Essays on the foundations of game theory* (1990) you will find a section where he says that, unfortunately, we get into a kind of impasse. We get this infinite regress linked to the common knowledge problem. For example, I drive frequently from Aix to Marseille. You have the *autoroute* and parallel to it is the *route nationale.* Say there is, one day, congestion on the *autoroute* and nobody on the *nationale.* I think: "Tomorrow I will take the *nationale.* But, wait a minute, these other drivers are intelligent too, so they will take the *nationale* tomorrow, I would do better to stay over here. But, wait a minute, these drivers are pretty intelligent so they can make that step too…" It is actually not logically possible to reason to the solution of these kinds of problems that people are supposed to be solving in game theory.

You can surely define an equilibrium, and say that if we were there nobody would want to move. But then you get to the problem of how we get to this equilibrium—the exact same problem that we have with general equilibrium.

One way out is evolutionary game theory, which does not have people reasoning. You simply identify individuals with strategies, and strategies that do better reproduce more, while strategies that are doing worse disappear. That is extremely mechanical; it drops any reasoning on the part of individuals and, as I said earlier, it relies on too simplistic an interpretation of evolution.

For certain specific, local problems, game theory is a very nice way of thinking about how people might try to solve them, but as soon as you are dealing with a general problem like an economy or a market, I think it is difficult to believe that there is full strategic interaction going on. It is just asking too much of people. Game theory imposes a huge amount of abstract reasoning on the part of people—far more than in standard economics where you only need to know the prices and your own preferences.

That is why I think game theory, as an approach to large scale interaction, is probably not the right way to go. But I still think that a really important insight comes out of game theory: as soon as people start to worry about the fact that what they do has an impact on what other people do (and they start to think about it), that makes life very different.

You favour agent-based modelling as an alternative method to study agents' interactions. Can you sketch the characteristics of this method?

There are two possible approaches to agent-based modelling. The first approach is to start with a very simple, rudimentary model that can be solved analytically. Then you generalize it and simulate this more general model. We know the analytical results in the simple model and the question is whether these results continue to hold in the less restrictive model. You find such an approach in the chapter on fish markets in my new book (Kirman 2011, chapter 3). In the simple model with two sellers and many more restrictive assumptions, we worked out analytically whether people increase the probability of going to the seller from whom they made the most profit in the past. Then we ask what would happen if we tried to generalize the model. Since the results can no longer be derived analytically, we simulated what happens.

The alternative approach—the artificial life approach—gives people pretty much arbitrary rules to start with, and lets them choose different rules as they go along, and then you see if anything emerges from that. That was the Santa Fe stock market approach (e.g., Arthur et al. 1997): throw these individuals into the pot and then you look at the soup and see if anything has happened.

An objection to this second approach is that you have so much freedom. You can choose completely arbitrarily the very basic rules that you give people to start with. David Colander at some point raised this objection. He said that, if you are reasonably clever, you can just give the right rules to get anything you want to come out.

So a legitimate objection to a lot of agent-based modelling nowadays is that the specification is often not justified; one just puts down rules which seem intuitive.

So for you it is essential to have some sort of empirical justification?

I am not asking for extremely tight justifications—either empirical or theoretical—but you need some. It seems to be reasonable to assume that people stick to rules which seemed to have worked well in the past. If you start to put in more specific rules in your model, I would want some empirical justification, ideally. And again, I would come back to the same idea: watch economic agents.

For instance, when the Soviet Union collapsed, all these kiosks sprang up in Moscow around the metro stations. You had all these people who started trading for the first time. We went to this new market and collected data over three months about the prices they charged and we asked them: "How did you choose the rule you use? How did you decide how to change the price?" The answers were very interesting. Some sellers would say: "Well, I look at the going price in Moscow". Others said: "I check the price I paid for the articles when I got them, and then I add a mark-up", or "I try to match the average of what is around me", or "I try to beat the lowest price around me". So you had all these different rules and then we ran little simulations and found out that you could have higher prices occurring at one metro station compared to another, not because the income levels of this area was higher or lower, but just because of the way these rules interact. We never published that, but it was really interesting.

What would you say are the most important insights delivered by agent-based modelling in economics so far? Whose work do you particularly admire?

Well, firstly, I do not regard myself as an agent-based modeller. I find it interesting and I think it is a good thing to do, but I would not define myself as a practitioner of that art. But let us start with people like Bob Axtell and Josh Epstein. In their book *Growing artificial societies* (Epstein and Axtell 1996) one finds a lovely illustration of how these methods can be used to explain things that we would find very difficult to explain with standard models. They developed what they called Sugarscape—a computational model inhabited by these little people who just wander around eating sugar. They are programmed so that they seek out where there is more sugar, but of course because they all end up going to, and consuming in, the same places, there is less sugar there, and so they move on. The society organizes itself and you can see cycles happen: more sugar grows here, they eat it and then they all move on, and later they come back when sugar grows there again. You get these cycles and patterns emerging which you could not have generated with an ordinary sort of model.

So we have some nice models which develop these sorts of insights to do with the emergence of collective self-organization. But I do not think that we have seen many market models, many economic models, which have such a striking result. For the moment we somehow have not made the leap. We have a few good examples though, such as the Schelling models (Schelling 1978). What was Schelling's motive? It was to show you that what happens at the macro level is very different from what happens at the individual level. So for me, that was an early example of agent-based modelling, and a beautiful one. As you can see from the beginning of my book, I am a complete admirer of Schelling.

Leigh Tesfatsion is someone who has made an enormous effort to gather people together around agent-based modelling (e.g., Tesfatsion and Judd 2006). So she should get a lot of credit for pushing this field. Blake LeBaron is another—an economist at Brandeis who works on financial markets. He has done some very nice agent-based models in which he manages to reproduce phenomena which are quite difficult to produce with standard models for stock price series (e.g., LeBaron, et al. 1999). You could also regard some of the work that Buzz Brock and Cars Hommes did as being agent-based. Again, those are people

who really develop interesting models; they have theoretical foundations but they also use simulations.

I am not a great fan of people who put together a model and then just turn the handle. Those who just say: "Now I will inspect the whole parameter space, to see what can happen and what does not happen". I do not think this is insightful. But there are some people who really have a vision. To come back to emergence, they find that they have generated interesting emergent phenomena from simple agent-based models.

Some scholars seem to believe that agent-based modelling offers a powerful tool for policy analysis. For instance, Robert Axtell says: "With that machine full of agents you can find out what processes give the empirical data on the distribution of real income and wealth. Then you can move them out for three years and see what happens. I think it's a new way to do macroeconomics; it's a new way to do policy" (interviewed in Colander, et al. 2004, 285). What kinds of policy questions do you believe agent-based modelling is capable of addressing?

When I gave a talk at the Bank of England, one of the directors, Spencer Dale, said: "I love what you do, I really think it is interesting, but you know, I have to make monetary policy. What am I going to say when we are sitting down and dealing with day to day policy matters? It is so much easier with regular models because we know what we are supposed to say and we know what people expect us to say and we can talk around that. But if you want me to use what you are saying, then you really have to explain to me how I transit".

I think that we are not actually ready for that yet. Of course, my view is that orthodox macroeconomics is not ready for that either, it is just that we are used to doing it that way.

There are some cases where you can start to see how these models can be used to guide policy though. I was asked by Bridget Rosewell, who is now the Chief Economic Advisor of the London Council in the UK, to help her look at what would happen if there was a pollution permits market. One way of doing that was to try and set up a simple theoretical model. But the alternative was to work out how people would behave in such a market; consider what strategies they might use, and then simulate that and see what comes out of it. With Nick Vriend, we found out that by changing the rules that the agents use we could

change the volatility of the resulting prices a lot. That was what was worrying the government at the time, i.e., that these prices were going to be too volatile. By looking at the agent-based model, we found out that we could reduce the volatility substantially by not allowing people to do certain things.

So there are already some areas where these models are being used, but to say that we should move on to doing macroeconomic policy now on the basis of agent-based modelling is too much. There are some optimists like Joe Stiglitz, Mauro Gallegati, and others who are starting to make macroeconomic recommendations on this basis. Yet other people are not so convinced.

What about the work of Andy Haldane of the Bank of England? He seems to be trying to use these sorts of techniques to try and get a better understanding of what "Too Big to Fail" actually means—at the moment it seems a bit of an empty slogan?

That is exactly the sort of thing he has in mind. And you are right that at the moment "Too Big to Fail" is just a slogan. We hear that if this bank went down, the others will go down too, so we should not allow the first to go down. But that is not articulated in the same rigorous way that you would require of somebody doing mathematical economics.

The question is, at what point does this interdependence become so important in the network that you cannot let one of our nodes go down? We can look at a banking network and consider if letting this node go down will take the whole system down with it. The step we have to take now is not only to look at the size of these nodes but to ask, what are the links? What are the arrangements that exist between these banks? What happens if one bank does actually get into trouble? Are the contractual arrangements such that the overseeing authority would see that the whole system would automatically get into trouble or are they such that the authority could just stay to one side? If we do not have the answer to that sort of question, then we do not know whether the trouble will propagate. So I think the next step is to really study the markets closely, and that is what central banks can do because they have the data—they know to what extent Barclays is tied up with HSBC, and so forth. Then, once we know more about what these links are, we can connect them as nodes and links. We have to study these links and nodes before we can say if this network is likely to collapse.

Summing up then, something like system resilience would be a policy dimension that network theory and agent-based modelling can contribute to.

Yes, and it is the resilience of the network that matters more than the resilience of individual banks. We used to say, "let us put these restrictions on individual banks", but if you only worry about individual banks, you do not know what "Too Big to Fail" means, because "Too Big to Fail" is really about the impact on others. It does not mean *what happens to you*, it means *what happens to the others*. Putting restrictions on individuals just means we do not want anybody to fail.

Continuing with policy matters, let us return to the problem raised by David Colander that you mentioned earlier. He expresses some scepticism about the capacity of agent-based modelling to offer clear policy advice: "When you have models with multiple equilibria, path dependency, nonlinear dynamics, endogenous tastes, institutional restrictions, and hysteresis, there are so many degrees of freedom that theory presents little in the way of restrictions on policy. With that many degrees of freedom, a sufficiently capable modeler can devise a theoretical model to support any policy" (Colander 2003, 20-21). What is your opinion about the prospects for agent-based policy analysis? Do you think that we can gather enough evidence to constrain our model specifications?

In the end the people building these models should be able to see which of the assumptions are driving the results. In general it is not true that all the assumptions have equal weight. Some might be changed without much impact. So in the end we should start to see which assumptions are driving a particular result.

Let me give you one very simple example. A lot of people have been trying to reproduce the exponential distributions or Pareto-type distributions of all sorts of phenomena, whether it be stock-market prices, the size of firms, city size, or whatever. This distribution appears all over the place; it is amazing. Now if you look carefully at what is going on here, it turns out to be always based on the following observation: if you take the case of cities, what you do is to say that when someone wants to move to a city, the probability to do so is proportional to the size of the existing city. With that rule, you automatically produce a Pareto distribution of city sizes. The same thinking is behind macro models of capital distribution. What happens

is that firms that have the most capital are those who will get more capital.

We see these wonderful Pareto distributions all over the place, yet they are in fact due to a very simple underlying mechanism. I think the same thing will happen elsewhere. It is not true that just because there are so many of these things we will never see the light; we will start to see which assumptions are driving the results. And once we do that it will make our job much easier.

In your recent book, one finds a lot of terms like 'understand(ing)' and 'explain(ing)'. In contrast, you seldom use terms like 'predict(ing)', 'forecast(ing)', 'control(ing)', 'intervening' and 'policy making'. What would you reply to someone who believes that economics is primarily in the business of predictions and policy recommendations and that, while your enterprise is perhaps valuable for explanatory purposes, it is of little use in the pursuit of these goals?

My wife says the same thing to me. She says: "Whenever you talk, I always have the feeling of somebody who is looking at an ant nest or a beehive and is very interested in what is going on in there and is really curious about it, but is not particularly worried about making it work better. In some sense, you are a curious observer rather than someone who is actually in the business of doing something". I think that is a legitimate criticism. I do not know whether that is my nature or what. I just find these things very interesting. I reason a little bit like an entomologist.

So, on forecasting, if you believe in this sort of systems approach, forecasting is a very difficult exercise. Just look around now at people forecasting, and people have these big, very sophisticated models. But, when you look at the discussions about what the growth rate in the European Union will be, or the growth rate in the United States, say, you see how quickly these things are revised. From one month to the next, the French government says: "Well, we have come down from 2.5 to 1.7". Is that not a big change? It is, in fact, a very big change; it makes a huge difference in terms of what we had better do with the deficit, and so forth.

I think we will do much better by looking at the nature of the evolution rather than trying to predict "this is going to happen"; trying to say, would this type of change that is happening lead to a more positive evolution or a more negative evolution? That is something we

can probably say something about, but saying that the growth of GDP will be 1.2, 2.3, or whatever, I just think we are not in that game.

In using this stuff to actually make policy recommendations, I think one has to be pretty modest. In fact, I am not sure how much it is really about using economics and how much it is a matter of having a vision of the world. When Hugo Sonnenschein said to me that macroeconomics was about wisdom, well, I have come to believe more and more that he was probably right. In some sense, macroeconomics is a lot about experience and rather little about formal analysis. But that is just a personal view, and a bit of a cop-out—a way of saying: "Sorry dear boys, I cannot do this!"

We also want to ask you about your feelings toward the use of mathematics in economics. Many heterodox economists argue that there is too much mathematics expected of economists, and that the profession has become obsessed with being overly formal. How do you think economists should use mathematics?

I find it strange that we should worry about a tool, that this tool should somehow be a criterion for judging work or be the subject of criticism. Mathematics is just a way of simplifying a problem, perhaps wrongly at times. John Chipman did a survey of international trade theory at one point, which was published in *Econometrica* (Chipman 1965), and there he says that sometimes mathematics turns out to be useful because it enables you to frame things in a clearer way. For him, solving certain problems in economics without mathematics is a bit like crossing the Channel by swimming. It is an admirable feat and everybody applauds it, but it is probably not the easiest way to cross the Channel. So, in a sense, avoiding mathematics in principle does not seem to me to make any sense. But becoming obsessed with mathematics does not seem to make any sense either.

In fact, I do not think the real issue has to do with mathematics and non-mathematics. Mathematics is just a tool, and really whatever tool that you can find around, well, that is fine. But somehow there is now a hierarchy, and mathematics is thought to be a superior thing to do. Recently, in our group, they refused a PhD student that I wanted them to take because she did not have hardcore training in mathematics. And she had done courses in business and so forth and wanted to work on behavioural finance. But they said: "This is not a serious person". This in my opinion is a very poor criterion, because her making some

progress on this particular problem does not necessarily require her to be a mathematician. There are lots of people around, like Akerlof and Bob Shiller, who do not use very high-powered mathematics, but do have good insights.

Many of my colleagues think that Schelling should not have won the Nobel Prize. When I ask them why, they say, "there is almost no math in what he does!" And it is certainly true that he does not use difficult mathematics. But he has difficult and really interesting ideas. Why should I judge him on the mathematical tools he uses?

Years ago, I was involved in organising conferences with Christopher Zeeman, who was one of the founders of catastrophe theory and the head of the mathematics department at the University of Warwick. He used to organize *rencontres* between mathematicians and people from other disciplines, and we organized one between economists and mathematicians. We had some great mathematicians—John Milnor, Steve Smale, Rene Thom, and others—wonderful mathematicians. And on the other side, we had Gérard Debreu, Hugo Sonnenschein, Werner Hildenbrand, and a whole group of very distinguished mathematical economists. After the first two, three hours, I think it was Milnor who said: "We all know that you guys can do mathematics, you do not have to show us. Everybody does his own thing. You want to show us that you are good at doing certain sorts of mathematics; that is fine. But we are interested in the economic problems. We thought that you were going to tell us about economic problems and we were going to use our mathematical tools to help you. But all you are telling us is the mathematical tools that you use and how you are doing well with them. But that is not going to create much". I think that was absolutely right. After that, the economists were rather silenced and started shifting in their seats uncomfortably. Debreu never said very much anyway, but it was clear he was very insulted, because basically he liked to think of himself as a mathematician.

Is it possible for you to summarize your diagnosis of the state of macroeconomics?

I would say that macroeconomic theory has gone down a blind alley in the sense that we have locked onto a particular model: general equilibrium. But it is not really general equilibrium, I mean, it is a one-man model! In particular, it has become mathematically sophisticated without representing the fundamental features of the macro-economy.

So I would say that people like Kydland and Prescott, and so forth, people like that, deserve their Nobel Prizes because they changed the way that people do macroeconomics. But in my view it was not a positive change. I think we have gotten away from worrying about the macro-economy as a system with interdependence, and so on, and become obsessed with this particular vision of how it works. One predominant idea is that of external shocks—and in particular the idea that the shocks that happen to the economy should essentially be the technological shocks. As Joe Stiglitz said, what could we mean by a negative technological shock? That people forget what they could do before?

So we have this idea that we have a system which is in equilibrium and that every now and then it gets knocked off the equilibrium by 'a shock'. But shocks are part of the system! We have gone down a track that actually does not allow us to say much about the real, major movements in the macro-economy. In the end, we should be more interested not in the periods where the economy is running along relatively smoothly, but in the periods where it changes. People typically say: "Well, this is not a normal period, and we analyse what happens in normal periods, and all of this is about deviations from that". But we should be studying non-normal periods, instead of normal ones, because that is what causes real problems. And we do not do that.

So my vision of the state of macroeconomics is that it somehow has the wrong view: an equilibrium view and a stationary state view. But what is important and interesting about macroeconomics is precisely when those two things do not hold. How can you talk of equilibrium when we move from 5% unemployment to 10% unemployment? If you are in Chicago, you say "Well, those extra 5% have made the calculation that it was better for them to be out of work". But look at the reality; that is not what happens. People do not want to be out of work. It is a tragedy for these individuals; it affects their identities. It upsets me a lot to think that people just say: "Ah, only another 5%; we handled this rather well". Millions of people are out of work, and we are not worried about that?

That is the major failure in macroeconomics. It does not address the serious problems that we face when we get out of equilibrium. And we are out of equilibrium most of the time.

That was the diagnosis, so now what treatment would you propose? And how do you think the profession should reform?

Ah, the profession. Well, as Buzz Brock says, we should open our minds, but not so much that our brains fall out. I think that we should take on board all sorts of different approaches to macroeconomics and to looking at how markets function. We should start to incorporate empirical evidence, rather than getting obsessed with extremely limited models. We should try to keep thinking about all these things that impact on the economy and see whether we can incorporate some of them, and maybe drop some other ones. We should not be totally focused on producing a closed-form model that you can solve and then say that it is a representation of what we see out there. If you do not look out there, this model will always be detached from reality.

That is a criticism of the profession. If you want to succeed, you have to publish in good journals. What the good journals publish are basically advances on previous work. That is absolutely reasonable, and it is understandable that the profession should have a lot of inertia in it. But it also should not be so locked in that anything that is more innovative cannot get into these major journals. I think, if you look at the *American Economic Review*, it is actually not really so bad, because it does incorporate quite a lot of behavioural economics, experimental economics, and so forth. But in macroeconomics, I think, it is extremely conservative. If you produce a model which is not in line with what was being done before, it is very difficult to publish.

There was this young economist, I think he was at UCLA, who wrote to me when I wrote this paper called 'Whom or what does the representative agent represent?' (Kirman 1992). He said: "Dear professor, I really agree with what you said. I think that it is intellectually absolutely right. Unfortunately, I am a young macroeconomist who is an assistant professor. I build models based on a representative agent. I know how to do that, and I know how to publish that. And I need to get tenure. Once I have got tenure, maybe I will be able to turn around and start to think about the sort of models that do not use the representative agent, but unfortunately, what I think will happen is that by then I will have got into the habit of doing it. I will publish my articles, get a decent reputation, I will get promotion, and I will probably never think about this again. But anyway, thank you very much for the insight!"

That is a bit depressing!

No! I thought that was extremely honest. That is just the way it is. It is very difficult. Curiously enough, places like the *Journal of Political Economy* allow for Schelling-type models and so on every now and then. So it is not true that the profession is a solid block against anything innovative. But there is a natural suspicion of things that cannot be reduced to a standard equilibrium.

So what reforms are necessary? Well, the way that the profession reacts is to create new journals, right? For example, from the outset I was involved with the *Journal of Economic Behavior and Organization*, which was really considered a marginal, weird journal. But nowadays it is a very acceptable journal; it is well thought of. That is the way that these changes will happen. I would like to believe in pure natural selection, but I think that there is a lot of inertia in the profession.

Every now and then I get invited to conferences in neurobiology. There people are really interested by what they are doing. Once they start telling you about what they are doing, you know, you just cannot stop them. But in discussion at economics conferences, it is usually about who is going to get a promotion where, who published in this or that journal. We should change that.

If you go to an experimental economics conference, it is more like that. They are more excited about what they are doing. But if you go to a macroeconomics conference, all you want to do is get out! It is the price you have to pay. You have to be there. They are all there, listening very seriously to each other.

So ultimately, are you optimistic or pessimistic about the profession's future?

Well, as Keynes might have said, 'in the long-run' I think that things will get a lot better. I think that people will realize that economics is a wonderful and exciting subject and they will stop treating it as an analytical exercise which is independent of reality. A lot of people do empirical work of course, but not imaginative, exciting empirical work. Often it is rather routine. But I have more sympathy for the individual who gets down and starts to analyse some data for a particular market—the market for wheat or something—and really looks at how it works, than I do for the person who builds the n^{th}-generation DSGE model. I think that the first individual is trying to understand what is going on. The tools that he uses may not be super exciting but he adds

to our understanding of the economic world. That is what it is about. That is what I thought when I came to economics, I thought: "this is really about understanding how economic phenomena happen. What a wonderful, exciting subject".

REFERENCES

Arthur, W. Brian, John Holland, Blake LeBaron, Richard Palmer, and Paul Tayler. 1997. Asset pricing under endogenous expectations in an artificial stock market. In *The economy as an evolving complex system II*, eds. W. Brian Arthur, Steven N. Durlauf, and David A. Lane. Reading (MA): Addison-Wesley, 15-44.

Binmore, Kenneth G. 1990. *Essays on the foundations of game theory*. Cambridge (MA): Basil Blackwell.

Chipman, John S. 1965. A survey of the theory of international trade: part 1, the classical theory. *Econometrica*, 33 (3): 477-519.

Colander, David. 2003. Post Walrasian macro policy and the economics of muddling through. *International Journal of Political Economy*, 33 (2): 17-35.

Colander, David, Richard P. F. Holt, and J. Barkley Rosser, Jr. 2004. *The changing face of economics: conversations with cutting edge economists*. Ann Arbor: University of Michigan Press.

Churchland, Patricia S., and Terrence J. Sejnowski. 1992. *The computational brain*. Cambridge (MA): MIT Press.

Debreu, Gérard. 1974. Excess demand functions. *Journal of Mathematical Economics*, 1 (1): 15-21.

Epstein, Joshua M., and Robert L. Axtell. 1996. *Growing artificial societies: social science from the bottom up*. Cambridge (MA): MIT press.

Föllmer, Hans. 1974. Random economies with many interacting agents. *Journal of Mathematical Economics*, 1 (1): 51-62.

Goodwin, Richard M. 1951. The nonlinear accelerator and the persistence of business cycles. *Econometrica*, 19 (1): 1-17.

Gordon, Deborah M. 2010. *Ant encounters: interaction networks and colony behavior*. Princeton: Princeton University Press.

Härdle, Wolfgang, and Alan P. Kirman. 1995. Nonclassical demand : a model-free examination of price-quantity relations in the Marseille fish market. *Journal of Econometrics*, 67 (1): 227-257.

Hildenbrand, Werner. 1994. *Market demand: theory and empirical evidence*. Princeton: Princeton University Press.

Hildenbrand, Werner, and Alan P. Kirman. 1976. *Introduction to equilibrium analysis: variations on themes by Edgeworth and Walras*. Amsterdam: North Holland Publishing Co.

Hildenbrand, Werner, and Alan P. Kirman. 1988. *Equilibrium analysis: variations on themes by Edgeworth and Walras*. Amsterdam: North-Holland Publishing Co.

Kirman, Alan P. 1992. Whom or what does the representative individual represent? *The Journal of Economic Perspectives*, 6 (2): 117-136.

Kirman, Alan P. 2011. *Complex economics: individual and collective rationality*. London: Routledge.

Kirman, Alan P., and Miriam Teschl. 2010. Selfish or selfless? The role of empathy in economics. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365 (1538): 303 -317.

LeBaron, Blake, W. Brian Arthur, and Richard Palmer. 1999. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23 (9): 1487-1516.

Mantel, Rolf R. 1974. On the characterization of aggregate excess demand. *Journal of Economic Theory*, 7 (3): 348-353.

Morishima, Michio. 1984. The good and bad uses of mathematics. In *Economics in disarray*, eds. Peter John de la Fosse Wiles, and Guy Routh. Oxford: Basil Blackwell, 51-73.

O'Connor, Timothy, and Hong Yu Wong. 2009. Emergent properties. *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta.
http://plato.stanford.edu/archives/spr2009/entries/properties-emergent/

Oullier, Olivier, J. A. Scott Kelso, and Alan P. Kirman. 2008. Social neuroeconomics : a dynamical systems perspective. *Revue d'économie politique*, 118 (1): 51-62.

Schelling, Thomas C. 1978. *Micromotives and macrobehavior*. New York: W.W. Norton & Company, Inc.

Sonnenschein, Hugo. 1972. Market excess demand functions. *Econometrica*, 40 (3): 549-563

Tesfatsion, Leigh, and Kenneth L. Judd. 2006. *Handbook of computational economics: agent-based computational economics*. Amsterdam: North-Holland Publishing Co.

Vinković, Dejan, and Alan Kirman. 2006. A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103 (51): 19261-19265.

Alan Kirman's Webpage: http://www.vcharite.univ-mrs.fr/pp/kirman/