

The complexity vision in economics

by Duncan K. Foley

Notes prepared for the National University of Ireland Galway Workshop on Computable Economics, March 21--25, 2005.

Economic aggregation

A fundamental problem in economic theory is the relation between the actions of individuals, households and firms, which produce, consume, buy, and sell, and the social phenomena these actions generate, such as markets and prices. We apprehend economic reality largely through statistical measures, such as GDP, employment, and price indices, which aggregate individual actions and transactions.

Every economic theory must propose a viable conceptual and methodological relationship between these levels. Since what we observe are largely statistical aggregates, some theories propose relations directly between aggregates, such as aggregate consumption functions or Phillips' curves. Under what circumstances are we justified in reasoning in this way? Must valid relations among aggregates be rooted in a rigorous analysis of individual behavior (micro-foundations)?

Aggregation, explanation and prediction

One reason for economics' preoccupation with these issues is the common observation that statistical aggregates are much stabler and more regular than individual behavior, due to the effects of the Law of Large Numbers, and therefore more amenable to mathematical explanation. Furthermore, economists and the important consumers of the knowledge they produce, such as politicians and the educated public, are often more interested in aggregates such as economic growth rates, employment rates, and the like than in explanations of individual behavior.

Statistical self-organization in thermodynamic and economic systems

Thermodynamic systems tend to a statistical equilibrium with well-defined emergent properties such as pressure and temperature. The second law of thermodynamics asserts the strong stability of this type of statistical equilibrium (maximization of entropy) in thermodynamic systems. It is the unique definition of entropy in these systems that allows us to link their aggregate properties directly through a state function. When we know the volume and energy of a given fluid system, we also know its temperature and pressure. This was the first type of self-organization that physicists understood clearly.

Marginalist economics was an attempt to apply physical ideas of equilibrium to economic exchange and production. Economists expected to be able rigorously to produce the same kind of linkage between macro-aggregates that characterize classical thermodynamic systems. Expressions such as $MV = PT$ for the relations among the money supply, velocity of money, price level, and transactions reveal this underlying program.

Economic agents

The neoclassical or marginalist economic agent is defined by a well-behaved preference ordering representable by an ordinal utility function $u[x]$. The gradient of the utility function $u'[x]$ represents the *relative offer prices* at which the agent will trade commodities. Two agents with different relative offer prices can find a mutually advantageous voluntary exchange. An *equilibrium* of a system of such agents is an allocation of commodities at which no further voluntary exchanges are possible (the *Pareto set*). (Walrasian equilibria constitute a small subset of the Pareto set.) Neoclassical economic agents thus behave exactly like equilibrated thermodynamic subsystems interacting with each other.

The existence of at least one economic equilibrium is assured if there is an economic problem of scarcity. The stability of the economy toward the equilibrium set is robust given any mechanism of voluntary exchange. The equilibrium set, however, is generally large and indeterminate.

Quasi-linear economies in which all agents' utility (or profit) functions can be written as $u^j[x^j] = x_0^j + \bar{u}^j[\bar{x}^j]$ where x_0 is the same good (wealth) for all agents also have well-defined entropies $S = \sum_j \bar{u}^j[\bar{x}^j]$ maximized on their equilibrium set. In these economies the *marginal utility of wealth* is constant, and changes in holdings of wealth absorb all the income effects. The micro-foundation of Keynesian economics is a quasi-linear economy.

Indeterminacy of equilibrium in general economies




Figure 1

The movement from non-equilibrium endowments to equilibrium (the Pareto-set) in economies is irreversible, non-unique, and indeterminate.

Indeterminacy of equilibrium in quasi-linear economies




Figure 2

In quasi-linear systems, however, the prices and allocation of the non-linear goods are invariant in the equilibrium set, and this degeneracy makes these properties of the economy determinate.

Path-dependence in complex systems

We can see in the case of general economies where the marginal utility of wealth changes along the path from non-equilibrium endowments to the equilibrium set that the final allocation and prices of all goods are indeterminate. It was to address this problem that Walras introduced the fiction of the auctioneer and Edgeworth the device of recontracting, in the hope (which turned out to be wrong) that the Walrasian set of equilibria reachable by recontracting at which the value of agents' final bundles is equal to the value of their initial endowments would be unique and therefore determinate.

In fact, economists stumbled on the phenomenon of *path-dependence*, which caught the attention of most physicists only much later. The reason that simple thermodynamic systems have statistically self-organized equations of state is that they have no memory of the irreversible path through which they reach equilibrium.

Physicists attack this problem using the methods of *statistical mechanics*, the combinatorial analysis of the micro-states of a thermodynamic system consistent with its macro parameters such as energy and volume. This method is not available to study the micro-structure of economic agents.

Walras' attempt to define unique economic equilibrium

Walrasian economics attempts to resolve the problems of the indeterminacy of the robust and simple definition of equilibrium by arbitrarily choosing one path of exchanges from an initial endowment to the contract set. The Walrasian idea is that an equilibrium (in the sense above) is characterized by particular relative prices (common to all the agents), and to use these same prices to define the path of exchanges leading to equilibrium. This leads to a backward logical situation, since we need to assume the answer to the problem (which equilibrium the economy will reach) in order to analyze the problem itself (what path will the economy follow from its initial endowment to equilibrium). Indeed, this logical backwardness in Walrasian reasoning led in the twentieth century to a tremendous mathematical complication of equilibrium logic.

In order to carry out the Walrasian program we must define a *Walrasian equilibrium* as an equilibrium whose prices give the same value to each agent's endowment and her final commodity bundle (that is, keep the agent on her budget constraint). If we knew these prices to start with, we could plausibly suppose that the economy would find this equilibrium through a set of exchanges at precisely the equilibrium price ratios. The question is how do we know what these prices are before the actual market process has taken place?

The question of the existence and stability of Walrasian equilibrium bristles with mathematical difficulties and paradoxes. In particular the question of finding a robust stability in equilibrium prices has remained elusive, and the issue of the existence of Walrasian equilibrium has been settled only by introducing powerful abstract mathematical principles into the argument which have no real economic foundation.

Non-equilibrium phenomena

Many important phenomena, including life itself, cannot be explained as features of entropy-maximizing thermodynamic equilibrium. From the thermodynamic point of view, these phenomena take place far from equilibrium. They are typically path-dependent, in that the further evolution of the system depends critically on the path it took to reach its current state, not just on the state itself. The analysis of these phenomena requires some method that can address their detailed dynamics.

Equilibrium, self-organization, and complexity

Classical thermodynamic systems relax to equilibria that are statistically self-organized and determinate. Their equations of state (entropies) represent the interaction of their macro-aggregates. The complexity program broadly speaking seeks to extend this type of analysis to path-dependent non-equilibrium systems by finding other forms of self-organization that are stable or quasi-stable in such systems. Complexity theory is in this sense a branch of thermodynamics, which studies the generalization of thermodynamic methods to path-dependent non-equilibrium systems. The influence of this methodological program on economics in the last sixty years has been immense.

« ‹ › »

12 of 22

John von Neumann and automaton theory

The key figure in the development of the thermodynamic program to embrace non-equilibrium path-dependent systems and to address economic problems was the mathematician John von Neumann. von Neumann's first attempt to address this problem, game theory in the pre-Nash form, gave a satisfactory determinacy in zero-sum games, but proved incapable of addressing non-zero sum games. von Neumann viewed Nash's "equilibrium" approach to non-zero sum games as unlikely to produce determinate results, and devoted the last years of his life to the study of *automata*, abstract representations of entities capable of complex interactions. Much of the foundation of complexity theory lies in the theory of automata. von Neumann's program was to derive emergent properties of interactive systems of automata from general properties of the automata themselves at various levels of sophistication.

« ‹ › »

13 of 22

Characterizing complex systems dynamically

Dynamical systems are characterized by their *stability* and corresponding *attractors*. The simplest class of dynamical systems have unique, asymptotically stable equilibria with point attractors. These systems tend to a unique state from any initial conditions. Conventional economics has been preoccupied with the effort to represent complex economic interactions in this form. The second simplest class of dynamical systems converge asymptotically to periodic motions with an attractor like a limit cycle. Much of Richard Goodwin's pathbreaking work addresses economic systems that behave in this way. The trajectories of both these types of system tend to converge dynamically over time, making it possible to predict their evolution with high precision from imperfectly measured initial conditions.

A third class of dynamical systems are *chaotic systems*, in which trajectories diverge over time. Chaotic systems have *strange attractors*, fractal self-similar sets towards which they converge, but on which they diverge. The dynamic divergence of chaotic trajectories makes it impossible to predict their evolution precisely from imperfectly measured initial conditions, but it is possible to characterize their long-run behavior statistically. While it is difficult to predict the exact position of a chaotic system on its attractor at a future point in time, it is often possible to predict the amount of time the system will be observed in any portion of its attractor quite accurately. Classical thermodynamic systems such as gases are chaotic dynamical systems in this sense.

Adaptive, self-organizing systems far from equilibrium

Complex systems that can maintain themselves far from dynamical or thermodynamic equilibrium in apparent defiance of the Second Law of Thermodynamics, such as the living cell, ecologies, the brain, and capitalist economies, cannot, viewed as dynamical systems, be members of any of these three basic classes. In stable dynamical systems structures disappear by being compressed toward the equilibrium attractor, while in chaotic dynamical systems structures disappear by exploding into all parts of the attractor. Complex systems must lie on the boundary between stable and chaotic systems, so that structures in their initial conditions can evolve without being destroyed. Like chaotic systems, complex systems may be stable in some dimensions (and chaotic in others) giving rise to *self-organization* either of a deterministic or statistical variety. It is possible to understand the self-organizing dimensions of complex dynamical systems without assuming that their overall trajectories are deterministically or statistically predictable.

Commodity economies are a good example. There are in the short run highly stable aspects to commodity economies, such as the emergence of prices (due to short-run entropy maximization), the equalization of profit rates, and the like. There are also stable and quasi-stable evolving structures (institutions) whose exact trajectory cannot be predicted. It is much more likely that we could predict the existence of asset markets and speculative formation of asset prices a hundred years in the future than what technologies will be or what the relative size of different sectors of an economy will be.

Cellular automata

Cellular automata are a particularly simple class of automata that live on a graph or network. At least in the simplest cases, all the nodes of the lattice are identical "cells" which can take on a finite range of "states". To begin with I will consider cellular automata that live on a one-dimensional lattice (the integer numbers, or the integer numbers modulo some number n , which makes the integer number line effectively into a circle by making cell n the immediate lefthand neighbor of cell 1. But there are some well-known higher-dimensional cellular automata, such as John Conway's Game of Life. The Game of Life is a two-dimensional lattice in which each node is a cell that can be in one of two states, "alive" or "dead". As time passes (usually in cellular automata in discrete fashion, like the ticking of a clock) a dead cell comes to life if it is surrounded by exactly 2 or 3 living cells, and a living cell continues to live only if it is surrounded by exactly 2 living cells. It has been shown that there is a particular starting configuration for the Game of Life that is a universal computer, in that it will carry out any computation a general Turing machine can.


⏪ ⏩ ⏴ ⏵

16 of 22


Representation of cellular automata

One-dimensional cellular automata are defined by the number of states, k , which Wolfram often calls "colors", and the radius or range of the neighborhood, r , which determines which neighboring cells affect each cell. The simplest cases is when $k = 2$, and $r = 1$, so that there are only two states, and each cell changes according the state of its immediate neighbors and itself.

We can visualize the evolution of a 1-dimensional cellular automaton by graphing its states (using colors) across a number of pixels, and then moving down as time moves forward. For example the Rule 30 cellular automaton (with two states and a range of 1) produces this type of output on random initial conditions.



Rule 90 produces this type of output:





Complexity Classes in Cellular Automata

Rule 30 produces a dense, apparently constantly changing, expanding pattern. Rule 90, on the other hand, produces a highly regular, repetitive pattern. These are examples of Class 3 and Class 2 cellular automata, respectively. Class 2 cellular automata produce regular, periodic patterns in both space and time that are equivalent to periodic or limit-cycle motions of dynamical systems. Class 3 cellular automata produce highly mixed, constantly varying output and are equivalent to chaotic dynamical systems. Class 1 cellular automata rapidly converge to a single state both in time and space, and are equivalent to asymptotically stable dynamical systems, as the following example (Rule 128) shows:



Class 4 cellular automata, however, exhibit quite a different type of behavior. It turns out that $k = 2$, $r = 1$ cellular automata are incapable of Class 4 behavior, so the "simplest" example has $k = 3$, $r = 1$:





There are several remarkable and thought-provoking features of this Class 4 cellular automaton. First, unlike Class 1 and 2 CAs, Class 4 does not settle down into a single state or repetitive pattern. On the other hand, the structures generated by a Class 4 system are not explosive like those of Class 3, and, in particular, sometimes highly elaborate structures collapse into much smaller ones. Class 4 CAs can maintain structures for long periods of time, and allow them to interact. Thus the Class 4 cellular automaton, though it is an extremely simple construct, provides us with a very abstract model of complex interactive systems.

Computation and complex systems

Complex systems have several important computational properties: *computational irreducibility*, *computational replicability*, *computational universality*, and *computational tractability*.



19 of 22

Computational irreducibility of complex systems


It is impossible to simulate the behavior of a complex system with a system less complex than itself. This implies that the only way to study the evolution of a complex system is to allow it, or an equivalent system, to evolve over time.



20 of 22

Computational replicability of complex systems


Many complex systems, such as cellular automata, however, can be generated by the application of simple rules of component behavior and interaction. Thus the evolution of a complex system which is exactly represented in an algorithmic model can be reproduced exactly in different runs of a computer program. The problems associated with understanding the behavior of complex systems are not, as a result, fundamentally problems of *approximation*.



21 of 22

Computational universality of complex systems

There is strong evidence (for example, from Wolfram's work on cellular automata and related systems) that systems as complex as Type 4 cellular automata are all capable of acting as Turing machines, that is, as universal computers. Thus any complex system, viewed algorithmically, can, in principle, emulate any other complex system (though different representations of the same system might involve immensely varying transparency and time-efficiency of the programs involved). This property of complex systems is particularly important for the economic representation of *expectations*, since human beings interacting in an economy are themselves complex systems making an effort to understand and predict the behavior and interaction of other individuals, also complex systems.



22 of 22

Computational tractability of complex systems

The main tool for the exploration of the properties of complex systems is *computer simulation*. With this methodology come a huge advantage and an important vulnerability.

The advantage of representing complex systems such as economies explicitly as simulation programs in computers is that there is essentially no limit on the type of behavior that can be represented. Analytical methods in economics often force researchers to make simplifying assumptions (such as convexity of production sets, reducibility of economies to representative agent form and the like) to assure tractable closed-form representations of outcomes. Simulation programs are not limited in this way, since adding a line of code to represent an exception or addition to agent behavior in certain circumstances is straightforward.

The vulnerability of simulation methods lies in the difficulty of showing how general and representative the results of particular simulations are for a broad class of models. If we can demonstrate business-cycle type fluctuations in the interaction of investing firms represented by one algorithm, for instance, how can we be confident that similar interactions will produce similar cyclical fluctuations in similar, but not identical model scenarios?