Hayek Enriched by Complexity Enriched by Hayek

Robert L. Axtell[∇]

Department of Computational Social Science
Krasnow Institute for Advanced Study
George Mason University, Fairfax, Virginia USA

Santa Fe Institute, Santa Fe, New Mexico USA

Complexity Economics Programme
INET@Oxford, Mathematical Institute and Oxford Martin School
Eagle House
University of Oxford, Oxford UK

Abstract

Certain elements of Hayek's work are prominent precursors to the modern field of complex adaptive systems, including his ideas on spontaneous order, his focus on market processes, his contrast between designing and gardening, and his own framing of complex systems. Conceptually, he was well ahead of his time, prescient in his formulation of novel ways to think about economies and societies. Technically, the fact that he did not mathematically formalize most of the notions he developed makes his insights hard to incorporate unambiguously into models. However, because so much of his work is divorced from the simplistic models proffered by early mathematical economics, it stands as fertile ground for complex systems researchers today. Austrian economists can create a progressive research program by building models of these Hayekian ideas, and thereby gain traction within the economics profession. I suggest that instead of mathematical models the suite of techniques and tools known as agent-based computing seems particularly well-suited to addressing traditional Austrian topics like money, business cycles, coordination, market processes, and so on, while staying faithful to the methodological individualism and bottom up perspective that underpin the entire school of thought.

Version 1.21: 21 October 2014

∇ Prepared for the conference "40 Years After the Nobel: F.A. Hayek and Political Economy as a Progressive Research Program," Mercatus Center, George Mason University, October 2-5, 2014. I am grateful to Professors Alan Kirman and Nick Vriend for conversations in the Spring of 2014 on the topic of this essay, in Aix en Provence and London. Neither is responsible for errors of omission or commission. ***Draft: comments welcome!***

I Perspective: Simple Models vs Complex Reality

"[T]he safety that comes from acknowledging only the facts one likes is fundamentally incompatible with science. Sooner or later it must be swept away by the forces of history." Robert Laughlin and David Pines (2000)

For almost two decades the similarities between the emerging sciences of complexity and Austrian economics have been progressively pointed out (Vaughn 1999; Vriend 2002; Koppl 2009; Rosser 2009; Barbieri 2013). Specifically, F.A. Hayek's research is often highlighted as being particularly close to current topics at the frontier of complexity (Caldwell 2004; Gaus 2006). In this article I will situate some of Hayek's ideas within the complex systems perspective, eschewing exegesis wherever possible in favor of models and movies. It will be demonstrated that certain of Hayek's ideas are special cases of broader conceptualizations and are therefore enriched by complex systems ideas. I will also claim that other aspects of Hayek's work represent a 'fresh' interpretation of current thinking about complex adaptive systems and present opportunities for complexity researchers to learn something new. But first, some background.

What is *complexity*? It is a scientific approach that has been around for a long time (Fermi, Pasta and Ulam 1955; Lorenz 1963) but which has come into its own in the last forty years—from the mid 1970s (e.g., Feigenbaum 1978; Mandelbrot 1983)—driven by advances in nonlinear mathematics (e.g., Guckenheimer and Holmes 1983), on one hand, and the exponential growth in computing technology on the other. It has given us a new vocabulary—chaos, fractal, sensitive dependence on initial conditions, bifurcation, strange attractor, self-reproducing automata, artificial life—and new meaning to old terms like *emergence* (Morowitz 2002) and *evolution* (Beinhocker 2005). It teaches that the natural world can display exotic dynamics, that predicting the future can be impossible *even if* the equations governing the future are known *exactly* (!), and that the biological and social worlds, made up as they are of heterogeneous, living, cognitive things, are more complex than the physical world. This approach to science is *complementary* to conventional approaches, providing insights into certain kinds of systems that were previously not accessible to quantitative science, like the weather, traffic jams, and financial market dynamics. But at the same time it can also be viewed as a *substitute* for previous methodologies, particularly when those are too simple or crude vis-à-vis real-world complexity. In *complex systems*, aided by modern computation, the distributed, decentralized interactions of disparate objects are analyzed and the ways in which patterns at the level of the whole system emerge from such interactions are studied. Fluid mechanical turbulence (Frisch 1995) , super-conductivity (Bardeen, Cooper and Schrieffer 1957) and synchronization phenomena (Strogatz 2003) are three familiar examples from the natural sciences. In *complex adaptive systems* the component parts have the ability to change their behavioral repertoire in response to their environment or each other. Quorum sensing (Fuqua, Winans and Greenberg 1994) and immune response (Perelson 1988) are microbiological examples, while social insect behavior (Seeley 1995; Bonabeau, Dorigo and Theraulaz 1999; Strogatz 2003) and animal herding (Gueron, Levin and Rubenstein 1993; Gueron and Levin 1993) are instances from ecology.

When I think of Nobel Prize-winning economists who have made major contributions to the modern science of *complex systems*, Kenneth Arrow comes immediately to mind for his roles in (a) helping found the Santa Fe Institute (SFI) some 30 years ago, (b) bringing Brian Arthur to SFI where the latter has been a core member of the faculty, (c) creating a dialog between economists and physicists (Anderson, Arrow and Pines 1988), (d) inviting many well-known economists to visit SFI, and so on. I also think of Herb Simon because of (a) his influential MIT lectures in the 1960s, including "The Architecture of Complexity" (Simon 1996), (b) his early work on empirical and mathematical aspects of power laws (Ijiri and Simon 1964; 1967; Simon 1955; Simon and Bonini 1958; Ijiri and Simon 1977) which we understand today as a kind of 'signature' of complexity, (c) his myriad efforts in artificial intelligence (AI), which would set the stage for *distributed AI* (DAI) and eventually *multi-agent systems* (MAS) computer science, and (d) his efforts to model cognition using complex systems ideas. Tom Schelling's role is prominent due to (a) his pioneering work with what is today called agent-based computing in the guise of the model of residential segregation that bears his name (Schelling 1971; 1971; Casti 1994), (b) his little gem of a book *Micromotives and Macrobehavior* which anticipated many of the themes that would later animate researchers at SFI, including the El Farol model (Arthur 1991; Bell and Sethares 2001; Bell, Sethares and Bucklew 2003) and the closely related 'minority game' (Johnson, Jefferies and Hui 2003) which owe much to Schelling's model of seminar attendance, (c) his 'standing ovation' model which has been a touchstone for complex adaptive systems thinking (Miller and Page 2007), and (d) his discussion of traffic dynamics which foreshadowed both actual traffic models with emergent dynamics (Nagel and Rasmussen 1994; Nagel and Paczuski 1995; Nagel and Barrett 1997; Nagel, Beckman and Barrett 1998) and general models of distributed, decentralized systems (Resnick 1994).

But predating all of these canonical contributions by almost a generation is the work of Hayek, who was an early proponent of the 'bottom up' perspective, broadly-construed, that is today at the heart of the complex systems worldview. In what follows I shall compare certain ideas from the burgeoning complex systems literature with analogs from Hayek's work, to investigate the extent to which Hayek anticipated modern developments, to assess whether his ideas are still relevant, and to determine what, if anything, the modern conceptions owe to him. I will consider more than a dozen areas of overlap between emerging ideas about complex systems and Hayek's writings. For each I will summarize Hayek's conception, compare and contrast newer thinking from complexity, and conclude with an assessment of how things might develop in the near term, including how Austrian economists might advance and extend the current state-of-the-art. My selection of topics is in no way exhaustive of the intersection between complexity and Austrianism. Immediately clear in these comparisons is that, while Hayek deserves significant credit for articulating many important ideas relevant to the social and economic sciences, his methodology was unconventional (by today's standards), with his work looking more like philosophy than modern economics: he uses little mathematics, the work is 'thick' descriptively but at a qualitative level without systematic use of quantitative data, yet he purports to develop general results. I conclude wondering whether, taking Hayek's writings on the whole, he would have wanted to be known as a scientist as opposed to a moral philosopher? I argue that his work can provide a suitable basis for a scientifically-grounded Austrian economics, but this would seem to require adoption of methods, largely computational in nature, that have been little leveraged by Austrians to date.

0 Simple Systems vs Complex Systems (and Pathologies of Interpreting the Latter in Terms of the Former)

Methodological individualism, with its clear Austrian origins (Menger 1985 [1883]), has been a hallmark of the economics profession for a very long time. It manifests itself in subfields in different ways. For example, in the theory of general equilibrium each individual consumer is, in principle, represented—its preferences, its endowments—yet these individuals do not interact with other individuals directly but rather with economy-wide phenomena like price vectors, interest rates, and so on. In the theory of the firm it is the norm to consider each business entity as a unitary actor. It is *as if* every company, no matter how big or small, was operated from the top down by a single mind, perhaps its founder. Another way of saying this is that the multi-agent character of firms is generally neglected. In game theory agents interact directly with one another but the only 'solutions' studied are ones in which each agent has no incentive to deviate, i.e., agent-level equilibrium. Such solutions are *simple* in the sense that they do not permit more complex notions, such as each agent having a suite of strategy functions and the composition of all the agents' functions constituting a strategic ecology.

To this kind of limited individualism economics, as a social science, has to add a behavioral model. Rationality is one such behavioral specification, yet it is also an ideal type. It is not empirically grounded in actual human behavior (Newell and Simon 1972) but is used rather because it is mathematically simple, i.e., analytically tractable (Simon 1957). For at least 60 years psychologists have known, from experiments, that people are not rational (Simon 1956). For at least this long the importance for economics of treating human cognition as bounded has been forcefully advocated (Simon 1955).

While methodological individualism is clearly an important starting point for the creation of a modern economics, it comes with a cost. For the real world is so variegated that it sometimes appears that every individual is *sui generis* and every state of the world idiosyncratic. However, in *most* economic models we have only one or at most a few *types* of agents, because to include more is technically (analytically) challenging. That is, in trying to be loyal to the methodologically individualist stance one encounters severe mathematical difficulties in representing the true heterogeneity of people and so it is often expedient to simplify one's analysis by having just a few agent types or perhaps a single representative agent.

Yet societies are multi-level, with activities at the individual or agent level potentially quite different from those at the macro-social level. In figure 1 these distinct levels are depicted as state x and y, respectively. We imagine that x is an element of a very high dimensional space, representing as it does the composite state space of all the people in society. The vector y is considered to be much smaller in dimension, perhaps some sufficient statistics concerning the society, its population, economy, and so on.

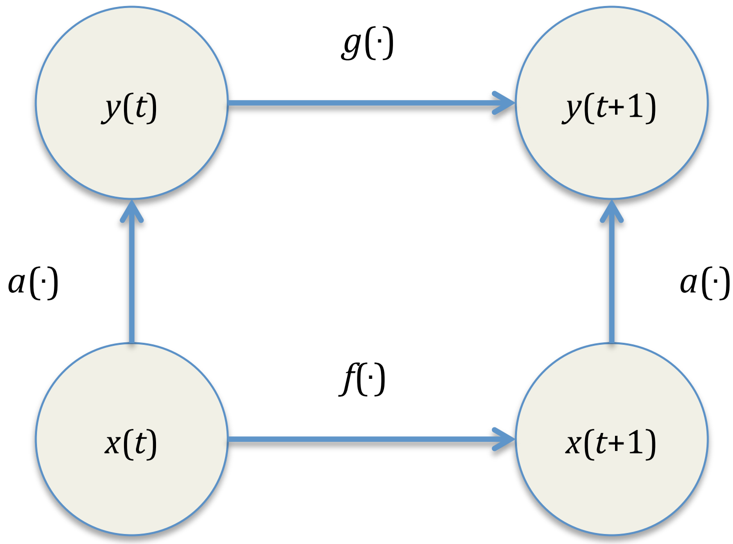

Figure 1: Conceptual description of a multi-level system, with the high dimensional bottom level being the agent level while the upper plane is the aggregate or macroeconomic level.

In this figure the macrostate, $y(t)$ is derivable directly from the microstate, $x(t)$, according to $y(t) = a(x(t))$, where $a(\cdot)$ is an aggregation operator, which does not depend on time. The micro-dynamics are assumed to be known and it is possible to march the microstates forward over time as $x(t+1) = f(x(t))$ while the macrodynamics, $y(t+1) = g(y(t))$ are somehow derived analytically, statistically or computationally. There are vast literatures studying the properties of aggregation functions and the resulting macrodynamics given the microdynamics, in economics (Ando, Fischer and Simon 1963; Gollier and Zeckhauser 2005) and in other fields in which the distinction between micro-behavior and macro-phenomena manifests itself (Epstein 1983; Luckyanov, Svirezhev and Voronkova 1983; Iwasa, Andreasen and Levin 1987; Gard 1989; Iwasa, Levin and Andreasen 1989; Gueron, Levin and Rubenstein 1993; Nilsson and Görnerup 2008). Unfortunately, aggregation in this way is generally impossible: there either will not exist functions that perfectly aggregate microscopic states or else such functions will not be unique. To get around these difficulties economists have made large use of representative agents, but such abstractions bring their own problems.

It is a tenet of elementary philosophy that it can be hazardous to infer properties of one level from knowledge of properties at the other level. That is, the *fallacy of composition* starts out with the known properties of *x* and wrongly attributes them to *y*, while the dual *fallacy of division* occurs when the structure of *y* is used to infer properties of *x*. (The latter is known as the *ecological inference* problem is statistics and econometrics.) When the two levels possess distinct properties one or more of them may be *emergent*. For example, when the macro level is in a steady-state (equilibrium) condition but the micro-level is not we say that equilibrium *emerges* from the interactions of the parts. These ideas will be discussed at greater length below, so suffice it to say here that the multi-level character of economies is a thorny issue that can make the entire subject of economics seem riddled with paradox.

The purpose of the discussion so far is to suggest that conventional economic models are very simple, perhaps too simple: noninteracting agents in place of interaction, rationality substituted for realistic behavior, and the micro and macro levels viewed as only different in *scale* via the representative agent, not *qualitatively* different. I summarize these distinctions between the simple and the complex in economics in table 1, adding additional contrasts.

Economic conception	**Simple**	**Complex**
Quantity of agents	representative (one, few)	many (possibly full-scale)
Diversity of agents	homogeneous	heterogeneous (or types)
Agent goals, objectives	static, scalar-valued utility	evolving, other-regarding
Agent behavior	rational, maximizing, brittle	purposive, adaptive, biased
Learning	individual, fictitious play	empirically-grounded, social
Information	centralized, maybe uncertain	distributed, tacit
Interaction topology	equal probability, well-mixed	social networks
Markets	WMAD, single price vector	decentralized, local prices
Firms and institutions	absent or unitary actors	multi-agent groups
Governance	benevolent social planner	self-governance, emergent
Temporal structure	static, impulse tests, 1-shot	dynamic, full transient paths
Source of dynamism	exogenous, outside economy	endogenous to the economy
Solution concepts	equilibrium at agent level	macro steady-state (stationarity)
Multi-level character	neglected, dual fallacies	intrinsic, macro-level emerges
Methodology	deductive, mathematical	abductive, computational
Ontology	representative agent, *max U*	ecology of interacting agents
Policy stance	designed from the top down	evolved from the bottom up

Table 1: Contrast between simplistic conceptions of economic processes and a more realistic view

Taking the RHS of table 1 seriously means adopting a complex systems point of view. In terms of how to build models that have the features of the RHS I shall eventually argue that we have today a methodology for dealing formally with the 'Complex' entries of this table, so-called agent-based computing. So it would seem that progress could be immediate and forthcoming. However, the trick is that we cannot simply turn ON everything on the RHS, all at once, or else we are put to sea not knowing where we are or having sufficient navigational apparatus to get ourselves back to known ports. So much work to date, both mathematical and computational, has focused on taking baby steps from the 'Simple' column, usually one dimension at a time.

At this point it may be useful to many readers to unpack what is meant today by 'complex systems' in general and 'complex adaptive systems' in particular. I will first take an example from the physical sciences, fluid mechanical turbulence, to illustrate the former and one in which human behavior plays a large role, traffic flow, to depict the latter.

Fluid flow from simple to complex

Consider the flow of a viscous fluid, like water or motor oil or air, moving from left to right in a pipe of radius *R*, driven by higher pressure on the left than on the right, as shown in figure 2, a lengthwise cross-section of the pipe.

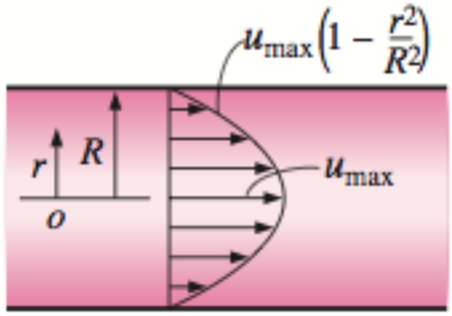

Figure 2: The laminar velocity distribution, $u(r)$, of a viscous fluid flowing in a circular pipe is parabolic, with the fluid in the center flowing faster than the fluid nearer the walls.

The local fluid velocity is parabolic, as indicated in the figure, depending on the radial distance, *r*, from the centerline of the pipe. Fluid in the center of the pipe moves faster than fluid closer to the boundaries. This simple flow, with each annular 'shell' of fluid proceeding down the pipe at its own velocity, interacting little with the fluid in neighboring shells, is called *laminar*. The parabolic velocity distribution is an *exact* result, derivable from the so-called Navier-Stokes equations (Bird, Stewart and Lightfoot 1960) and the equation of continuity. The maximum velocity, u_{max}, depends on *R*, the pipe length, *L*, the fluid viscosity, μ, and the pressure drop, ΔP, and is given by

$$u_{max} = \frac{R^2}{4\mu}\frac{|\Delta P|}{L}$$

where ΔP is measured as $P_{out} - P_{in}$ and is therefore negative, thus the absolute value sign. The overall flowrate, *Q*, is proportional to the average velocity $u = u_{max}/2$ times the cross-sectional area of the pipe, $A = \pi R^2$, i.e., $Q = uA$ and so

$$Q = \frac{\pi R^4}{8\mu}\frac{|\Delta P|}{L}$$

an expression known since the middle of the 19th C as the Hagen-Poiseuille law. By changing the length of the pipe or altering the pressure on either end of it the flowrate changes quantitatively, but the velocity distribution remains parabolic. For a very large pipe, like a water main or an oil pipeline, or a much smaller one, like a vein in your body, laminar flow and thus the parabolic velocity distribution well-describes steady fluid flow. (Oil and blood can have viscosities that are not constant but that is immaterial for the present discussion.)

However, laminar flow is not the whole story. It is a simple story, and arguably less interesting than what else can happen. Given a particular pipe (i.e., *R* and *L* specified) and a distinct fluid (so that μ is fixed), if sufficiently large pressure is applied to a laminarly flowing fluid a *qualitative* change in the flow can be observed. Increasing the pressure drop increases the average and maximum velocities of the fluid and thus its overall flowrate, and beyond a certain point the laminar flow regime is destabilized. In place of the steady streamlines of figure 1 the flow self-organizes into vortices and eddys that are essentially transient—they are dynamic and ever-changing. This situation is depicted in figure 3, a snapshot of an actual flow regime in a pipe having clear walls. Looking a moment later would reveal a different pattern.

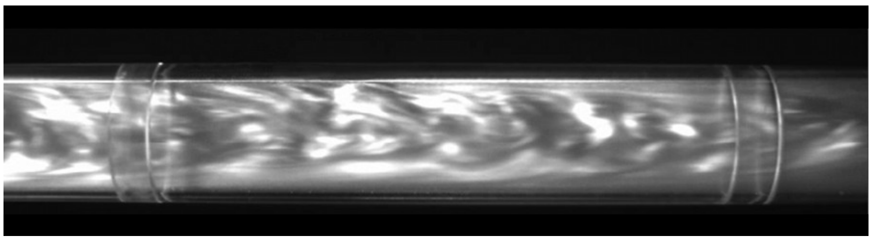

Figure 3: Photograph of turbulent flow of a viscous fluid in a circular pipe with clear walls.

Clearly it would be a mistake to imagine that the flow in figure 3 is well described by the velocity distribution of figure 2. In the turbulent flow there are a host of shapes and structures that have arisen spontaneously. These quickly evolve into other forms and patterns that are both spatially and temporally transient. These structural features of the flow field are constantly rearranging themselves and are more-or-less short-lived. The velocity distribution is now erratic, not parabolic as in figure 2. The fluid streamlines—paths traced out by individual fluid particles—were straight lines in laminar flow but are clearly irregular now, perhaps even chaotic. In laminar flow neighboring 'shells' of fluid slid by one another, resisted only by the fluid's viscosity. But in these turbulent motions the fluid is also mixing and dissipating significant energy, meaning it takes more power to push the fluid this fast—there is a price to be paid for creating all these internal structures, it turns out, a price in units of energy!

So turbulent flow is clearly very complex in comparison to the laminar regime, and if one imagines trying to describe it using the simple expressions above then significant errors will be made. But that is effectively what one does when the empirical features of the real economy, as on the RHS of table 1, are interpreted using the 'simple' approach of the center column.

Turbulence is an example of a complex system: the fluid reconfigures itself (self-organizes) into meso-scale structures—larger than the molecules the fluid is made of but smaller than the size of the pipe. How does the fluid accomplish this? How does the material 'know' how to reorient into three dimensional, transient structures, and what can be determined about the size and shape of these structures? Perhaps surprisingly to people who have not studied fluid mechanics, the answers to these questions are "we do not know". We do not know how to mathematically represent the large irregular fluid structures shown in figure 3, so we certainly do not know how to analyze how they change over time, and therefore we lack the ability to derive the velocity of the fluid as a function of its radial position, although this can be measured and has the approximate shape shown in figure 4. As the flowrate increases large-scale structures give way to smaller and smaller ones until the onset of 'fully turbulent' flow, for which it is believed that the velocity distribution can be represented as a random variable (Frisch 1995).

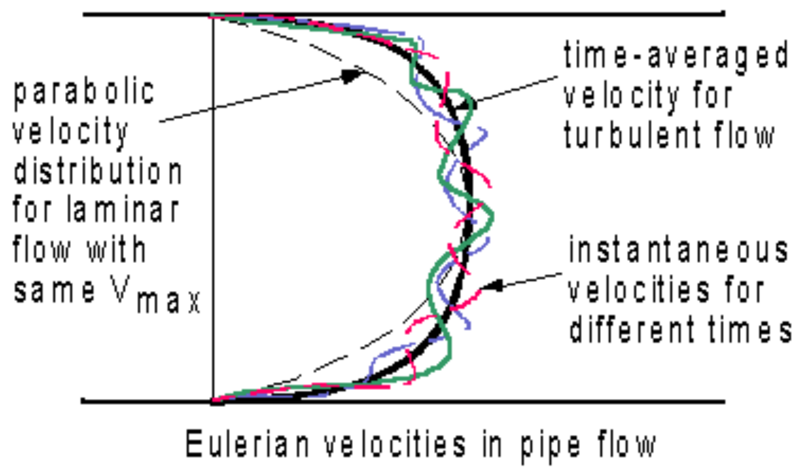

Figure 4: Approximate random velocity distribution, $u(r)$, in turbulent flow.

One way to find a 'signature' of a complex system is to look at its fluctuations. In laminar flow the velocities of individual fluid particles are not always prefectly parabolic, but rather vary to some extent. But because the individual 'layers' of fluid interact very little, these fluctuations are only weakly coupled to one another and are Gaussian in character, i.e., they are governed by the central limit theorem and represent 'mild 'randomness in the language of Mandelbrot (1997). However, once the flow becomes even mildly turbulent, velocity variations produced by the large-scale vortices and are non longer Gaussian, but are much more heavy-tailed, as shown in figure 5a, a log-log plot of the probability (vertical axis) of a velocity change of a specific size (horizontal axis). Heavy-tailed fluctuations are also characteristic of stock market returns on sub-monthly time scales (figure 5b).

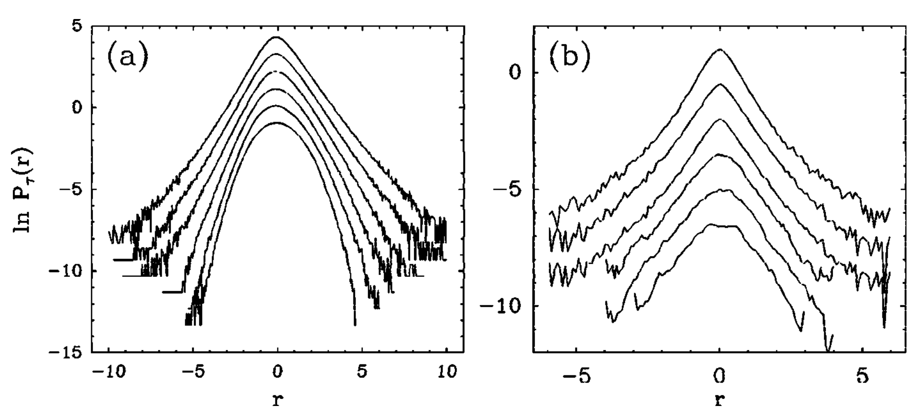

Figure 5: Heavy-tailed (a) velocity fluctuations in turbulent flow and (b) returns in financial markets.

Distributions of firm growth rates over time show a very comparable structure (Perline, Axtell and Teitelbaum 2006).

It turns out that the energy consumed by a fluid flow of this type, in which gravity plays no role (as long as the pipe is horizontal), is completely due to the shearing of neighboring fluid layers as they move past one another. The energy, E, put into the system by pumping essentially gets dissipated as internal energy (heat) due to the shearing actions, with most of this energy going into the fluid but some being transferred into the pipe walls. For laminar flow the rate of dissipation, $\dot{E}$, is given by

$$\dot{E} = Q\Delta P = \frac{8\mu L}{\pi R^4} Q^2$$

which is another *exact* result. From this expression, doubling the flowrate requires four times as much power. In turbulent flow there are different and more complex shearing forces due to the irregular motion of the fluid. Because so little is known analytically about turbulent flows, expressions for the rate of energy dissipation in turbulence are not available and we have to rely on experimental measurements. An approximate relation for the rate of energy dissipation for turbulent fluid flow in a smooth pipe is (Bird, Stewart and Lightfoot 1960)

$$\dot{E} \propto Q^{11/4}$$

From this expression it is possible to compute that doubling the turbulent flowrate requires about 6.7 times as much power! So a lot more energy has to be added to the system to get the fluid to flow turbulently. Where is all this additional energy going, how does it get transferred into the fluid and its surroundings, and how does this dissipation alter the flow? These are deep questions about which very little is known (Frisch 1995). Indeed, if you can answer them—and in this domain ‘answer’ means demonstrating that the Navier-Stokes equations have well-behaved solutions in the turbulent regime—then the Clay Institute for Mathematics will award you $1 million!

That turbulence is analytically intractable (at this time) is problematical. It means that we cannot say much about the interior of stars, since the plasma there convects in turbulent fashion. Nor can weather forecasters accurately predict where severe storms (hurricanes, tornados) will hit because the high winds associated with them are turbulent. Industrially, turbulence inside piping and fluid machinery severely limits the lifetimes and usefulness of such apparatus for commercial purposes. In shipbuilding, hull designers attempt to reduce fluid mechanical drag as far as possible, and some tricks are known for turning mildly turbulent flows into more laminar ones by releasing special polymer solutions into the paths boats will travel. But, in general, turbulence is a brick wall Nature has put up and which we do not know how to surmount analytically. It represents one of Nature’s great secrets that the human race has not (yet) figured out, mathematically!

But we do know a lot about the structure of turbulence from computational studies. By employing modern computation it is possible to *simulate* turbulent flows directly, either by numerically solving the Navier-Stokes equations using fine mesh representations of the flow domain (Hoffman and Johnson 2007), or else by computing the trajectories of individual fluid particles, a technique broadly-known as 'lattice gas' dynamics (Doolen 1990). Such computational approaches are capable of resolving large-scale eddy motion and do a reasonably good job reproducing complex flows, whether at the scale of individual pipes in a chemical plant or mass ejections from black holes in galaxies.

Professor Richard Wagner of George Mason has written that the distinction between simple and complex is like the difference between a parade and the people beside the road watching the parade (citation). The parade is organized, steadily flowing, scripted, more or less synchronous, even 'laminar' if the reader will permit some loose usage of terminology. The crowd is turbulent: some sit, some stand, some watch for a while and then head into the corner shop for a cup of coffee or a pack of chewing gum; some who live along the route watch out of their upstairs windows moving back and forth to the kitchen to refill their drink. The behavior of the crowd is irregular, unscripted, asynchronous. The crowd is not 'flowing' so it cannot be turbulent, but in its irregular undulations it pulses, its component people circulating.

A crowd is a complex *adaptive* system (CAS). Such systems have many features in common with complex systems, but are characterized by components that can learn, not merely changing their behavior as their local environment changes, but also having the ability to change their behavioral *rules*, whether through simple learning or outright innovation. Of course, social systems always feature adaptive individuals, so all manner of social forms and processes are some form of CAS, while physical and chemical systems—non-living, inert systems—are not adaptive in this sense. Biological systems, of which humans are an important example, of course, are CAS but many simple biological lifeforms look more like non-adaptive complex systems than CAS.

Traffic flow from simple to complex

As an example of a CAS not unlike the fluid flow system above, consider vehicular traffic on a single-lane highway. All vehicles wish to traverse the highway but if the road gets very crowded the total vehicular flowrate can decline. Driving rules represent the behavior of individuals. It is known that a wide-variety of 'collision free' rules produce the same kind of qualitative behavior and are capable of generating emergent traffic jams having properties that closely resemble real-world traffic jams, e.g., backward moving jam fronts. Figure 6 is a simple example of a few dozen vehicles moving around a circular road. In this model it turns out to be guaranteed from the driving 'rules' given the agents/vehicles that no collisions happen, and of course there is no 'road work' or other surface imperfections (e.g., potholes) present in this simple model to perturb the flow! So if we observe simulated traffic that is not perfectly smooth it is because of the *internal* behavior of the overall system, and not the result of *external* shocks.

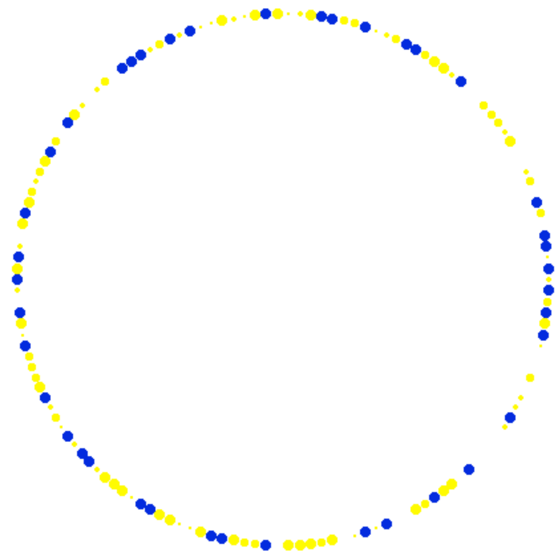

Figure 6: Snapshot from an agent-based model of single-lane traffic. (For a version of this model that has been unwrapped into a line visit the following website: https://ccl.northwestern.edu/netlogo/, download NetLogo, then open the “Traffic Basic” model under File>Models Library.)

When you run this model you see irregular flow of traffic with lots of starts and stops, accelerations, decelerations, and so on. When few vehicles are present the flow is highly regular, laminar-like. For a large number of vehicles the flow is slow and halting. But for intermediate numbers the flow is on the boundary between being smooth and jamming. It turns out that a plot of the total vehicular flow rate, in units of vehicles/time, is single-peaked. That is, at low vehicular densities average velocity is high but there are few vehicles on the road so throughput is low. At high densities the average velocity is low causing flowrate to again be low. But at intermediate densities the average velocity is reasonably high and the total number of vehicles is not small, so the total flowrate is high. It turns out that this peak flowrate happens quite near the point at which jamming starts and individual vehicle velocities become irregular. Stated differently, fast flowing traffic is on the edge of chaos!

As in the case of fluid flow, mathematical equations that reproduce these irregular traffic patterns can be written down but closed-form analytical solutions are not known. Our understanding of the spontaneous formation of traffic ‘vortices’ and ‘eddies’, even in this one-dimensional case, is almost exclusively due to computational models, mainly agent-based ones. What is more, as we move from abstract models as shown in figure 6 to real-world traffic situations, agent-based computational approaches are well-suited to adding the kind of driving heterogeneity, environmental forcing, roadway disruption, and even law enforcement activities that are found in modern highway systems.

As intermediate between the abstract model of figure 6 and full-blown real-world conditions, it turns out be relatively easy to run driving experiments on a circular roadway of *using real vehicles*. This has been attempted recently in Japan, summarized in general interest scientific journals (*New Scientist*, 2010 and *Science*, 2010). Figure 7 is a snapshot from a movie of this kind of experiment. When you watch it—see the URL in the caption—you see approximately the same kind of behavior among real drivers that was on display in the computational model, and the same kind of aggregate outcomes: stop-and-go, irregular flow around the loop that does not abate with time, backward moving jam fronts, and so on.

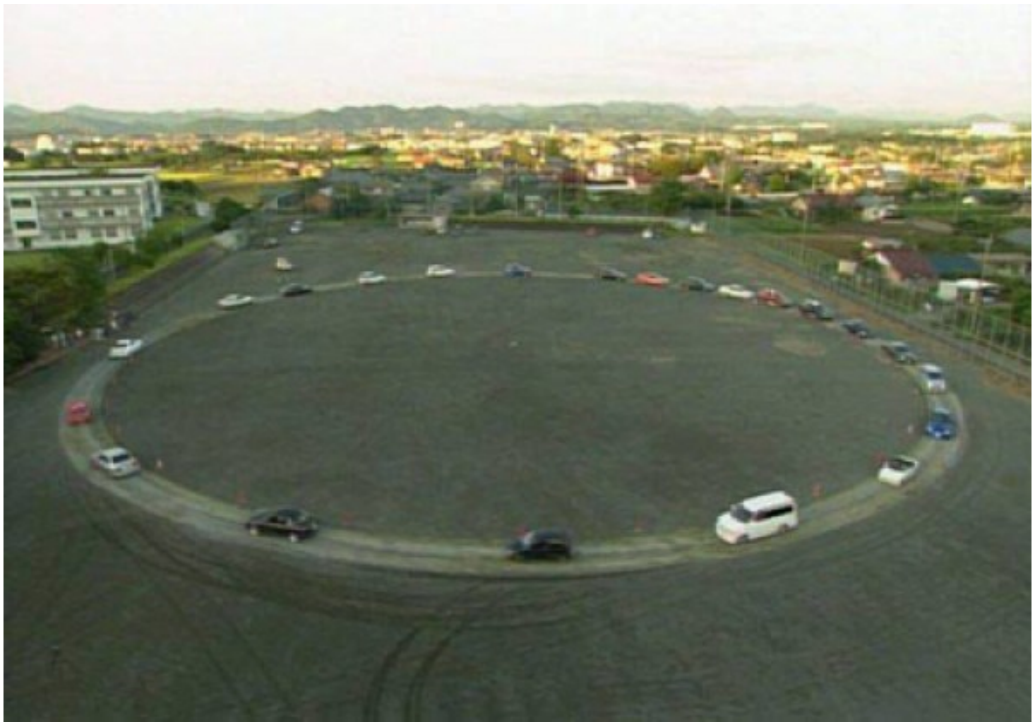

Figure 7: A traffic experiment in which irregular (stop and go) flow is observed despite there being no crashes, road work, potholes, bad weather, or exogenous events of any kind (https://www.youtube.com/watch?v=7wm-pZp_mi0).

While traffic phenomena of this type have been known empirically for nearly as long as there has been congestion, there have not been satisfactory ways to model it until recently. Indeed, until the advent of agent-based models of traffic some 20 years ago the norm was to treat traffic flow as if it were the flow of fluid in a pipe, resolved computationally. Such models do *not* exhibit the kinds of intermittent phenomena on display in figure 6 and 7. The metaphor of traffic as fluid flow is clearly too simplistic to address the kinds of phenomena that are not only observed in practice but also explicitly model-able using complexity tools. It is also the case that explanations of jamming phenomena as due to exogenous events are not necessary.

Hayek clearly anticipated many of the themes present in the modern science of complex systems, and particularly in CAS, as will be described below. Indeed, in "The Theory of Complex Phenomena" he uses very modern terminology (Hayek 1967) and addresses many of the major topics (e.g., institutions, evolution) that are active areas of research today. Perhaps this is unsurprising given his well-known interactions with Herman Haken, whose concept of 'self-organization' within physical and chemical systems (Haken 1987) was an early thrust area of complexity and remains relevant to questions surrounding the origin of life (Eigen and Schuster

1979). Hayek also had interacted with Ilya Prigogine, a pioneer in in the theory of non-equilibrium systems. However, it seems that most of Hayek's contributions to complex systems took place before any formal interactions with people who are today seen as pioneers of the complexity perspective. This leads me to believe that Hayek was essentially auto-didactic on complex systems, thinking about complexity from a social science point of view in much the same way that CAS researchers do today. With this brief overview of complexity now in hand we turn to the relations between Hayek's thought and certain notions from the complexity literature.

1 Spontaneous Order vs Emergence

Hayek is perhaps best known for his elaboration of the concept of *spontaneous order* (Hayek 1945). That individuals, left to their own devices, will behave in ways that are mutually beneficial is an old idea, going back at least to the Scotch political economists of the 18th C (Sugden 1989). Hayek's notion enriches these by arguing that the kind of welfare produced from the bottom up by autonomous people would be inherently corrupted by top-down interventions. In his "The Results of Human Action but not of Human Design" (Hayek 1967) he builds on his earlier writing from the 30s and 40s.

In complex systems there is a notion quite similar to spontaneous order, but one not associated with welfare properties, the idea of *emergence*. Usage of the term emergence has a checkered history in philosophy but has taken on a well-defined meaning in complex systems (Holland 1998; Morowitz 1998). Specifically, in the context of multi-level systems (e.g., figure 1) it refers to a pattern or regularity that arises at one level but not at another. Typically, in the context of complex adaptive systems, like the vehicular traffic example described above, we consider the actions at the lower level to be autonomous (e.g., driving) and these actions give rise to higher order structures (e.g., traffic jamming). The phenomena at the higher level may have no analog at the lower level—for example, a single vehicle cannot itself be a traffic jam—in which case the emergent phenomenon is associated with novelty. Examples of emergent phenomena are many once you start looking for them, from herding behavior in stock markets to peloton formation in bicycle races to the spatial clustering of city growth, all involve purposive individuals pursuing their own ends leading to correlated behavior. In physics, chemistry and biology the term 'self-organization' is sometimes used instead of emergence, as when it is said that vortices in a turbulent fluid flow self-organize or when polymer (macro) molecules (e.g., proteins) self-assemble.

So far we have discussed emergence primarily from the micro to the macro. There can also be downward emergence. Agents can create macrostructures that then constrain agents at the micro-level. This is perfectly obvious, and is in fact the essence of Hayekian spontaneous order. A path forms from the accumulated actions of individuals, and then once the path exists it is preferable (by most) to use it. But here we see the first glint of difference between the two fertile concepts. Hayek essentially always associated spontaneous order with positive outcomes—order is a good. But emergence simply suggests that patterns can arise, with no implication that any specific pattern generates positive welfare. Can there be spontaneous disorder? Of course there can, whether riots in the streets that destroy property or herding in stock markets that lead to bubbles or 'extraordinary popular delusions and the madness of crowds' (Mackay 1841) to cite one well-known source on the matter.

So what appears different in the usage of the two terms, by complexity scientists and Austrian economists, respectively, is that the latter nominally associate some kind of welfare measure to the spontaneous order, whereas emergent phenomena do not necessarily carry any particular welfare property. For example, when it is said that a market is a spontaneous order the implication is that those who are participating in it are gaining welfare from doing so, for else why would they do it? But take the NASDAQ financial market for example. Before decimalization c 2001-2 there existed a relatively common behavior within the so-called Small Order Execution System (SOES, pronounced 'sews' as in "the tailor sews my clothes"), called 'SOES banditry,' in which two people would be acting as market makers, one on the 'bid' side of the market and the other on the 'ask' side, separated by one or more units of spread (1/8s and 1/16s of a dollar at that time), working to execute orders, keep some inventory, and so on, but then a (typically larger) broker would jump into the gap in the spread and essentially execute all the orders the two previous market-makers had worked so hard to 'tee' up. In principle this kind of behavior is acceptable, but it is considered bad behavior. In the language of emergence SOES banditry is a specific kind of behavior that emerges on the NASDAQ because of the trading rules. But because of its negative welfare connotations, at least for the foiled market makers, perhaps it might not be considered a spontaneous order by Austrians, but rather something closer to a spontaneous disorder!

In thinking about the kinds of social orders that can arise spontaneously we can well imagine that free persons will attempt to arrange their affairs in ways that are mutually beneficial. If in the population there is a relatively 'flat' distribution of status then perhaps no one individual would wield unusual influence or power. One way to insure such a flat distribution of power is to have a legal system in which all people are treated equally, and Hayek was explicit about the rule of law being a necessary condition for 'good' social orders to arise. Without such equality under the law all kinds of social orders might arise that would not provide high levels of welfare for the individuals operating within the social order. For example, consider slavery, in which some individuals (slaves) have inferior standing and are thus dominated by others (slave owners). Social systems based on slavery can, of course, be quite stable, and there is no reason not to call them social orders, even spontaneous orders, but clearly they have inferior welfare properties. A more subtle example is the 'Jim Crow' social order in the American South that grew up after Civil War had made slavery officially illegal and had formally given African-Americans equal protection under the law. This social order institutionalized racial segregation and survived up through the mid 1960s, overturned only partially by a grassroots movement against it, augmented importantly by top down interventions from the Federal government, often in the guise of National Guardsmen. That white majorities were able to limit the freedom and liberty of minority groups is easily understood today, and well illustrates that spontaneous orders need not have positive welfare properties. Hayek was primarily focused on a different question, to wit, the deleterious effect of government on *good* spontaneous orders. But the example of the American South gives clear indication that top down interventions can have positive effects in limiting emergent phenomena from producing poor welfare outcomes.

A complexity economics example that is similar in sprit is an agent model for the emergence of social classes (Axtell, Epstein and Young 2001). Here we literally have spontaneous disorder. The metaphor is Dr. Seuss's *Sneetches* (1961) in which a meaningless body feature serves to order the society into two groups. In the story, once technology is discovered that can alter this feature, there results a perpetual 'war' between the groups until, in the end, the only thing being accomplished by individuals in the society is change to their bodies to comport with ever-changing group norms. In the model agents in a population are paired at random to play the Nash demand game. In this game it can arise that a group of agents, for historical and largely random reasons, arrive in a poorer welfare outcome than another group. While it is a theorem that such a state of the population will not last forever, it was demonstrated computationally in the paper that the transition time can be quite long. It is also possible that this model yields much more 'fair' outcomes, it just depends. But the main point is that it can generate final states that are spontaneously either well-ordered or disordered.

I summarize these considerations in figure 8, which qualitatively shows that there can be both emergent and non-emergent macrostructures, that spontaneous orders and disorders are the emergent structures that have welfare properties—positive and negative, respectively—and that spontaneous orders are clearly a special case of the more general idea of emergence.

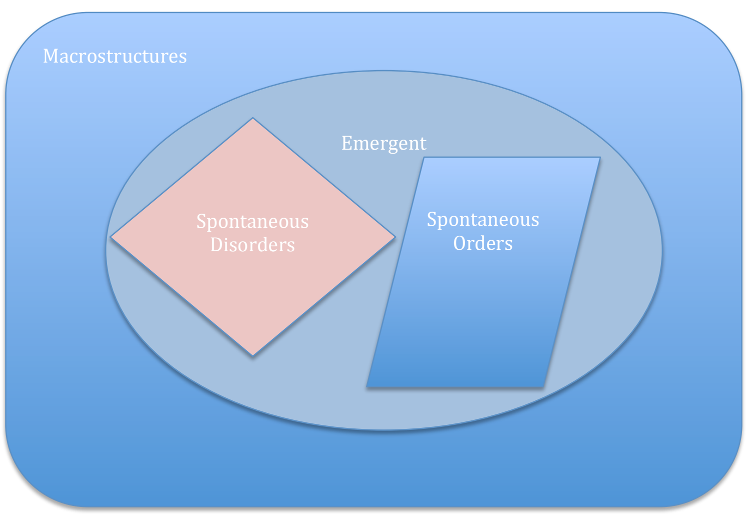

Figure 8: Spontaneous order as a special kind of emergence.

The point of this figure is not to diminish Hayek's contribution but rather to suggest that there is a broader conception of these ideas when one considers societies as multi-level.

To reinforce the idea that what *emerges* need not have a positive welfare property, perhaps it will be useful to consider a biological example, albeit a trivial one, not because it has any direct relevance to human social behavior, but rather because it clearly shows that the kinds of order produced from the bottom up can have negative welfare properties (Kauffman 1995). Imagine the self-organization of a malignant viral pathogen from benign components. If such a virus harms humans clearly there are negative welfare properties associated with its self-organization.

Two final points, perhaps of interest to readers who are not intimately familiar with the complex systems literature. First, emergence and self-organization represent large areas of study, and I note that I rarely see Hayek cited or the term spontaneous order used in those fields. Even when the term is used it is rare for it to be cited with its social science and Austrian origins. As an example of this I note that Steve Strogatz subtitled his excellent book *Sync* (Strogatz 2003) with *The Emerging Science of Spontaneous Order*, which is a perfectly informative way to describe the content of the book, with the last two words not intended to refer to Austrian economics (Strogatz, private communication). Finally, a point about scientific sensibilities. Sometimes people with similar training can come to quite different conclusions concerning the importance of certain ideas. While natural and social scientists are perhaps used to being surprised by Nature, it can be the case that engineers, medical doctors, and other professional do not look on 'surprises' as generative of positive welfare. In volume 1, number 1 of the *Journal of Autonomous Agents and Multi-Agent Systems*, which came out in the late 1990s and is today a thriving enterprise, the editors, all computer scientists but coming to their subject from quite different backgrounds, report in their opening essay about journal editorial policies, focus, administration, and so on, that there was one thing they could not agree on: emergence. To some of the editors the idea of emergence was an important topic that they wished to see discussed and highlighted in the research the journal published, while to the others, the computer engineers, they viewed emergent phenomena in computer systems as highly problematical, essentially unwanted 'side effects' that could wreak havoc on system performance, and should at all costs be avoided (Wooldridge 2002; Jennings, Sycara and Wooldridge 1998)! Clearly the welfare properties of emergent structures, if they can be assessed, matter in practice.

2 Tacit Knowledge vs Bottom Up Social Processes

Hayek, along with Michael Polyani (1974), articulated a theory of tacit knowledge of importance for economic production, distribution, pricing, and so on. For these writers the myriad social and institutional structures that were employed by people in economic activity were not designed, from the top down, by a single individual, but were rather the product of bottom up interactions among those individuals, with the resulting relations being only imperfectly codified and having origins that were not necessarily well understood by those who used them. The existence of large amounts of tacit knowledge (TK) in real economies is viewed as historically important insofar as it makes central planning an unworkable idea.

Given Hayek's sustained attack on Soviet and other collectivist economic planning activities it is no surprise that he made the role of TK a centerpiece of his analysis.

Before discussing modern developments in the theory of TK let me mention that over many years I have sat on a variety of funding agency panels to which work on TK has been proposed. Inevitably, such proposals begin by postulating the importance and prevalence of TK to some domain or other, often citing Hayek, Polanyi or both, followed by more or less formal statements of what TK is and how it can be modeled. Typically, there is enthusiasm on the panel for such work, initially, but once a proposal is discussed—when the proposer's *formal* ideas about *informal* knowledge have been understood and debated—the early support tends to evaporate, either partially or completely, often simply because reviewers come to believe that the kinds of things they think of as TK are not well-addressed in the proposal, or rather that the mere process of formalizing knowledge relations makes them no longer tacit. So while this observation—that it is hard to impossible to get research on TK funded—is not particularly strong evidence against the idea that TK is a meaningful idea, it does suggest that attempts to better understand it, especially to quantify it, are hard to accomplish in an unobjectionable way.

One way to conceive of TK—information about an economy that cannot be formalized, predicted, or understood *ex ante*—is to consider the *interactions* of individuals. It is more than a bit of a curiosity that neoclassical economics, the queen of the social sciences, conceived of as a science that deduces the aggregate consequences of individual actions, has historically not really permitted agents to interact directly with one another. Of course, in game theory the payoffs agents receive depend on how they play with or against others, but even there certain kinds of interaction, like preplay communication, are claimed to be irrelevant (i.e., cheap talk). In economics proper—think general equilibrium theory or financial market models—agents interact with price vectors and market institutions but not directly with one another. So let us consider models in which agents engage in explicitly social behavior through direct economic interactions. There is plenty of precedent for doing this. Back in the early 1960s, at Stanford and other places, a youthful crop of general equilibrium theorists relaxed some of the heroic assumptions of the Mackenzie-Arrow-Debreu model of general equilibrium (Mackenzie 1954; Arrow and Debreu 1954), in which no one trades until a notional auctioneer announces market-clearing prices, to produce market models in which prices changed or were heterogeneous across agents, all as a result of direct agent interactions (Negishi 1961; Uzawa 1962; Hahn 1962). However, one property of these models did much to interfere with their widespread acceptance by the mainstream of the profession: knowledge of the population's preferences and endowments was *not* sufficient to determine where the economy would end up. Rather, it was necessary to know which agents exchanged goods with each other in order to determine how the market would unfold. Stated differently, markets in which agents interact directly produce *indeterminate* final configurations unless the *exact* structure of the interactions was specified. The uncertainty and unpredictability of such models surely grated on neoclassical sensibilities and the reason why most researchers have not heard much about such models is because they never became a core part of the graduate curriculum. Similarly later, when

Radner and others introduced random variables into general equilibrium (Radner 1974), those models were never included in the core of Ph.D. education, for comparable reasons. All of which bring us back to TK. If it really matters, in actual economies, who trades with whom, which families eat at what restaurants, who buys a specific car, how individual wage bargains are struck, and so on, then can there really be any deterministic theory of general equilibrium? If where an economy ends up depends on the entire network of trading partners, and if this is not knowable *ex ante*, is this not a kind of TK? Where a small manufacturer may have extremely local information on which laborers she can employ and which products she will need to buy as inputs in order to produce a specific service or product, until such factors of production are actually engaged the exact trajectory of an economy will not be determined. Microeconomic factors such as these, across all firms in all industries, add up to an immense amount of finely detailed economic information that exists only in highly distributed records of decentralized processes that would be difficult to communicate to a central authority with any fidelity. Microeconomic minutiae of this type—indeed, it is so micro that it is almost nanoeconomic in character—matters and constitute a kind of TK about an economy, something easy to model with agents.

But how much does it matter? What if my dry cleaner gets their morning coffee from the convenience store on its left instead of the deli on its right? How much does this influence other, subsequent transactions? The neoclassical theory of general equilibrium gives us no leverage on this problem, because it assumes that all transactions occur at GE prices and denies that who trades with whom matters a wink. Agent-based computing gives us a way to characterize whether individual transactions matter or not. By building whole economies and making many realizations we can see how much, if at all, the ephemeral networks of who trades with whom matter. TK is real, and surely matters for the outcomes individuals receive, but does it have any impact on the long run configurations (e.g., growth rates) of economies, or is it always a second-order phenomenon? Bottom up models can help us determine this, but such models have not really been used in this way, yet. Such a research program would seem to be an important ingredient of any Austrian modeling program going forward, if Austrian ideas and models are to compete with neoclassical ones.

Lastly, let us consider another conceptual point about TK, one that manifests itself in the current research literature. It is conventional wisdom in macroeconomics that no individual actor is big enough to make a difference, in general, and that fluctuations in individual economic entities 'cancel out' at the aggregate level. So against assertions that TK matters a macroeconomist could argue that, while such considerations might matter for microeconomic outcomes, it is largely irrelevant at the macro level. Gabaix (2011) presents both theory and evidence on how microeconomic processes can matter for macroeconomic outcomes. Specifically, he studies the importance of micro actions to macro outcomes when firms are interconnected and have skew size distributions (Axtell 2001). He is able to prove that *aggregate* fluctuations will be heavy-tailed under such circumstances, independent of whether firm dynamics are heavy-tailed or not. In essence, the idiosyncratic shocks felt by individual firms amplify through interactions and lead to large macroeconomic fluctuations. He presents empirical evidence of the salience of these ideas for the U.S. economy, especially during the last few decades. If he is right it represents a whole new facet of TK that has heretofore not been much explored.

3 Market Processes vs General Equilibrium vs Market Processes

Having argued above that much of conventional, neoclassical economics is too simple to represent realistic economic processes, when it comes to the Walras-Mackenzie-Arrow-Debreu (Mackenzie 1954; Arrow and Debreu 1954; Debreu 1959) theory of markets (Weintraub 2011) it becomes incumbent to make the case with some definiteness. The WMAD model is called 'simple' in table 1 but it is by no means mathematically easy, relying as it does on nonconstructive fixed point theorems. Rather, it is simple in the sense that many of the social structures present in real economies, such as institutions, social norms, and so on, are abstracted away with the resulting model made highly parsimonious. I shall essentially argue that Hayek's conception of market processes does not comport with the WMAD theory and that, instead of viewing this as a *bug*, as would be the wont of the neoclassicals, it is reasonable to view it as a *feature*, since it provides fresh insight into the kinds of models of markets that we would like to build, if in fact we had the vocabulary to do so. I conclude by suggesting that a very Hayekian theory does exist but insofar as it is non-Walrasian it faces great hurdles, as do most alternative theories of markets, in the realm of scholarly opinion.

In Hayek (1945) a verbal description of market adjustment processes is presented without a formal model:

> It is worth contemplating for a moment a very simple and commonplace instance of the action of the price system to see what precisely it accomplishes. Assume that somewhere in the world a new opportunity for the use of some raw material, say tin, has arisen, or that one of the sources of supply of tin has been eliminated. It does not matter for our purpose and it is very significant that it does not matter which of these two causes has made tin more scarce. All that the users of tin need to know is that some of the tin they used to consume is now more profitably employed elsewhere, and that in consequence they must economize tin. There is no need for the great majority of them even to know where the more urgent need has arisen, or in favor of what other needs they ought to husband the supply. If only some of them know directly of the new demand, and switch resources over to it, and if the people who are aware of the new gap thus created in turn fill it from still other sources, the effect will rapidly spread throughout the whole economic system and in conséquence not only all the uses of tin, but also those of its substitutes and the substitutes of these substitutes, the supply of all the things made of tin, and their substitutes, and so on; and all this without the great majority of those instrumental in bringing about these substitutions knowing anything at all about the original cause of these changes. The whole acts as one market, not because any of its members survey the whole field, but because their limited individual fields of vision sufficiently overlap so that through many intermediaries the relevant information is communicated to all…The marvel is that... without an order being issued, without more than perhaps a handful of people knowing the cause, tens of thousands of

> people whose identity could not be ascertained by months of investigation, are made to use the material or its products more sparingly; i.e., they move in the right direction.

This passage was written well before the appearance of the WMAD model, and thus long before the Sonnenschein-Mantel-Debreu (SMD) critique of general equilibrium. From the perspective of WMAD, Hayek's description is of an adjustment process for a single good, with some reference to its relation to other goods as 'substitutes'. It is important to remember, however, that the WMAD theory does not permit *any* trading out-of-equilibrium, that is at non market-clearing prices, for this would create wealth effects, modify endowments, and so on. So Hayek's prose is really outside of the WMAD theory. But it is really not a model of the price level but of price adjustment? The trouble here is that under quite general conditions adjustment dynamics of the kind implied by the passage are not stable (Scarf 1960). They will *not* settle down, they will *not* yield fixed prices, in which case they are guaranteed to *not* be the *right* price? So this is an example of verbal reasoning that seems completely plausible but which is demonstrably wrong mathematically. However, one lesson from the SMD critique is that the WMAD theory is in some sense too general, for it places no constraints whatsoever on the relation between the microscopic structure of markets and the macroscopic properties of those markets. So why is WMAD the benchmark?

From Hayek's perspective WMAD is deeply flawed since it is single shot from endowments to equilibrium and does not permit any trading out of equilibrium. So why is not the passage above a *better* theory of market processes than WMAD. The problem is we do not have a simple mathematical theory of the kinds of things going on in Hayek's paragraph, at least nothing simple enough to teach first year graduate students. So we revert to WMAD!

There are many substantive critiques of WMAD, typically focusing on its unreality (Kirman 1989). Another class of criticism derives from the notion that any social process which is claimed to be 'bottom up' or an 'invisible hand' type process must be theorized/written in terms of individual actors and institutions that would realize the results claimed for the process (Nozick 1974; 1994). By any accounts, WMAD equilibrium fails on this score (Rothschild 1994). Indeed, as I have argued elsewhere (Axtell 2005), the WMAD theory is demonstrably *not* an 'invisible hand process' insofar as the only known ways to compute WMAD equilibria are so-called FNP-hard, among the most difficult problems in all of computer science. That is, the job of the Walrasian auctioneer becomes exponentially harder as the number of commodities increases. This is completely unrealistic and makes the WMAD model appear to be computationally naïve. Several writers have commented on the great power and elegance of the WMAD formalism, even going so far as to think of it as an invisible hand process (Stokey and Lucas 1989; Durlauf 2012), which surely it cannot be since it has no credible microfoundations.

If markets represent spontaneous orders, or more generally emergent phenomena, can we reconceive of them in Hayekian terms using the language and tools of complex adaptive systems? Can we build better models than WMAD, models that are computationally and behaviorally realistic? This is such a foundational research project, and could benefit significantly from Hayekian insights, that it would seem to be something that should be at the top of researchers' lists of priorities, whether those researchers are Austrian or not.

4 'More is Different' and 'Universality': Hayek vs Anderson

Hayek wrote importantly about the structure of complex systems generally and complex social systems more specifically. In "The Theory of Complex Phenomena" already mentioned he articulates a view that would become a touchstone for complex systems researchers a decade later:

> The "emergence" of "new" patterns as a result of the increase in the number of elements between which simple relations exist means that this larger structure as a whole will possess certain general or abstract features which will recur independent of the particular values of the individual data, so long as the general structure (as described, e.g., by an algebraic equation) is preserved.

The distinction between *qualitative* and *quantitative* features of a system are well-trod ground in science. A system with a given structure is normally believed to have certain quantitative properties that change with the numerical values of parameters and initial conditions, but fixed qualitative structure, determined by the overall form of the system. But in the above quotation Hayek appears to be saying something quite different and potentially very interesting, that the sheer scale of a system can determine what patterns emerge and suggests that new things can emerge as the system gets bigger and bigger. I assert that this astute observation was ahead of its time, and that in the social sciences it is today very much underappreciated.

By contrast a very well-known essay within the complex systems world, by physics Nobelist Philip Anderson, proclaimed that 'More is Different' (1972), and is now widely cited although it came long *after* Hayek. Anderson demonstrated that a wide array of physical phenomena can have the property that as their scale is changed (typically increased), the qualitative character of the phenomena that can be produced also change.

When Phil Anderson writes that 'more is different' in the world of physics he literally means that more particles, more electrons, more photons—in short, more of the same thing—can produce novelty. When it comes to Hayek writing "…'new' patterns as a result of the increase in the number of elements…" it may be that he had in mind people with heterogeneous skills, beliefs, desires, world views, and so on, in which case it would be unsurprising that some novelty were possible by simply increasing the scale. But because he conditions the increasing number to have the same "simple relations" I am prepared to believe that he is not talking about increasing the capabilities of the "elements" but rather simply their number. In fact, the second part of this long sentence seems to make clear that the heterogeneity of the "elements" is not what is at stake.

Independent of teasing out exactly what Hayek meant, there is another important concept on display here, one that is also little used in the social sciences but common in the physical sciences and important to the complexity sciences. When Hayek writes "…certain general or abstract features…will recur independent of the particular values of the individual data…" he is articulating what has come to be known as ‘universality’ in complex systems and certain branches of the natural sciences. Consider a fluid moving across a phase boundary, from liquid to solid say, by cooling. It turns out that the way this happens, macroscopically, is essentially independent of the microscopic structure of the molecules involved. That is, whether it is water or methane or carbon dioxide, the way the substance undergoes the freezing transition does not depend on the molecular structure of the chemicals in question. Such transitions are said to display ‘universality’ in the sense that the macrodynamics are common across a wide range of substances. Hayek is saying the same thing about the spontaneous orders that arise, that to some extent they do not depend on the idiosyncratic states of the elements they are composed of—those states do not matter for the transition process.

What we have here is the rarest of rare phenomena in intellectual history. There are countless examples in the social sciences where tools and techniques first applied in the natural sciences have been brought to bear on social data, processes, and systems. From new mathematics to network science to statistical techniques and computational tools, it has historically been a one-way street from natural science to social science. And with good reason, as the natural sciences are in many important ways *simpler* than the social ones, so progress, especially methodologically, almost surely happens first in the simpler domain. While most readers will not doubt the veracity of this perspective, I cite here a few volumes that are essentially attempts by practitioners of natural science to ‘colonize’ one or more fields of the social sciences (Montroll and Badger 1974; Gaylord and D'Andria 1998; Johnson, Jefferies and Hui 2003; Ball 2004; Buchanan 2007). It is unambiguous that Hayek was first on ‘more is different’ but Anderson is much more cited. Anderson does not cite Hayek, nor do social scientists Miller and Page in their textbook on complex adaptive systems (2007) in their discussion of ‘more is different’. In the same way that Weintraub may restore Mackenzie’s primacy to the Walrasian model, perhaps we can hope that future generations will recognize the social science origins of these two important ideas.

5 Flavors of Individualism: Atoms, Robots and Sociopaths vs Actors, AIs and Agents

The neoclassical model of consumer behavior, in which each person autonomously extremizes a scalar-valued utility function, is the workhorse of conventional theory. For every new economic environment the trick is to come up with a set of constraints and a utility function that when maximized in the context of the constraints yields comparative statics that comport with a researcher’s mental or empirical model of how the world works. Indeed, some researchers even believe their job is the “programming robot imitations of people” (Lucas to Klamer (1984)). In practice this model is well-known to be a poor description of humans from a behavioral perspective, with the list of cognitive biases growing almost daily. What is more, the idea of an individual acting without regard for other agents, or at least without direct interactions with any of them but only indirect interactions through constraints, is at least a peculiar model of humans, if not as my former Brookings colleague Henry Aaron has written, "a passable description of a sociopath" (Aaron 1994).

If people behaved in such ways the atomistic interpretation of methodological individualism would be valid. It would be as if people were robots[1]. Yes there are economic interactions that do not require much more than effectively anonymous interactions, as Buchanan has described certain market transactions (Bowles 1998). But other kinds of economic activity require deep and informationally intense interactions, not ones well-represented by autonomous optimization.

Game theory represents an advance over conventional microeconomics because it explicitly models direct interactions and its players may more reasonably be called *players*. Each has goals, strategies, and payoffs possible and realized. But Austrian economists are not major contributors to game theory proper, perhaps because they find the highly analytical focus of the field very narrow, the 'solution concepts' impoverished, and the static character of play very limiting, all things that were apparent to early mathematical economists like Herbert Simon in his 1957 review of the then new Luce and Raiffa text (Luce and Raiffa 1957).

In artificial intelligence (AI) it is common to build agents that take account of the state or actions of others in determining their actions (Russell and Norvig 2002). AIs with goals look more like agents, and the origin of multi-agent systems (MAS) computer science is largely an outgrowth of giving goals to distributed AIs.

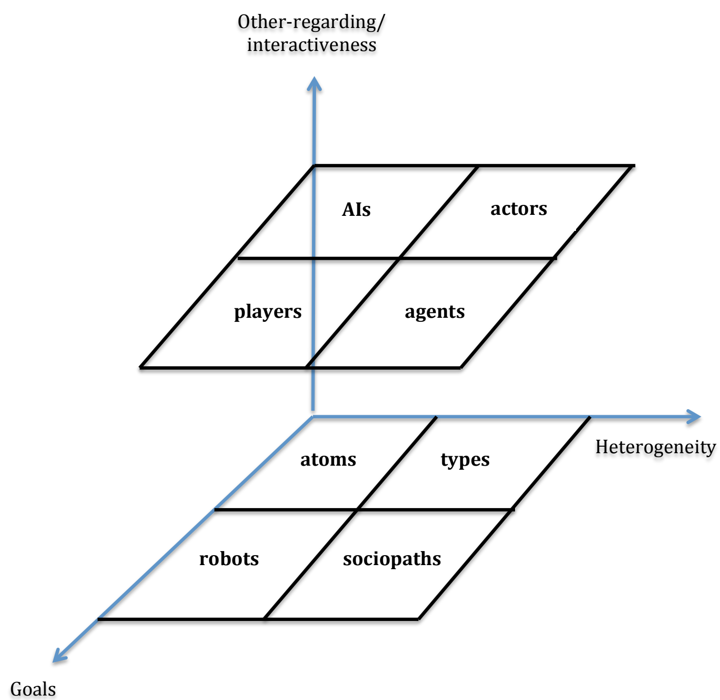

Figure 9: Distinct kinds of individuals in social science models.

[1] The word 'robot' derives from *robota* in Czech, meaning 'forced labor' in a 1920 play *R.U.R.*, an early 20th C dystopia.

Mitch Waldrop's book (Waldrop 1992) introduced many people to the field of complexity. In his recent update to that volume, to which he has added an afterword, he asks what the long-lived contributions to the SFI vision have been and what ideas have not withstood the test of time. For Waldrop the dominant invention has been agent computing.

6 Behavioral Rules and Realized Behavior: Hayek vs Cog Sci

In "Rules, Perception and Intelligibility" and then "Notes on the Evolution of Systems of Rules of Conduct" (Hayek 1967) the nature of human behavior is discussed from a pattern recognition perspective. That is, he discusses microstructures of cognition that go below the conscious level and into the perceptual. Contrast this with work that was going on roughly contemporaneously in the budding field of cognitive science. Here the founding fathers of this new approach to psychology were running experiments with human subjects to determine the extent to which 'production systems' could be made to reproduce laboratory observations (Newell and Simon 1972).

Over a decade ago Nicholas Vriend asked how Hayek might have employed rule-based reasoning systems (Vriend 2002). He created a so-called classifier system, based on the work or John Holland (Holland 1976), and coded it to reflect the kinds of behaviors relevant to information contagion processes he was studying. He clearly demonstrated that rule-based systems are a plausible alternative to optimizing, *max*{*U*} type specifications.

Today we know a lot about how people behave in specific environments than was known 50 years ago (Camerer 2003). A large number of cognitive biases have been observed, documented, and reproduced (Kahneman and Tversky 1979; Tversky and Kahneman 1981; Slovic, Fischoff and Lichtenstein 1982; Tversky and Kahneman 1986; Kahneman, Knetsch and Thaler 1986; Tversky and Kahneman 1992). That people use heuristics is undeniable (Gigerenzer, Todd and Group 1999; Gigerenzer 2000; Gigerenzer and Selten 2001). That these heuristics are efficient is provable (Michalewicz and Fogel 2000). That people learn is something that usually needs to be represented (Roth and Erev 1995; Camerer, Ho and Chong 2002; 2004; Ho, Camerer and Chong 2004; Shoham, Powers and Grenager 2004). All of which would lead one to wonder, who still argues that people are rational (Simon 1956)?

Unfortunately, the answer is that most of the articles published today in most economics journals *still* feature rational agents. Such agents have a rule, but it is more like a meta-rule from the perspective of behavioral rules: maximize utility. It takes a model to beat a model in economics today so the fact that the journals are yet full of maximizing specifications means that there are few alternative models of behavior. Rules of behavior, whether determined from laboratory experiments, random field trials, or back calculated from empirical data, represent an important alternative to *homo economicus* specifications. While the specification of behavior as optimizing essentially represents the normal *ontology* of neoclassical economics, and deviations from that norm are punished by referees and journal editors, it is also the case that at the boundaries of the profession other approaches are now on display.

For example, for about a decade, computer scientists working within the emerging *multi-agent systems* research thrust have been using all manner of game theory and optimizing agents in their models. But over the past five years or so there has emerged—partially as a response to the incredibly difficult optimization problems that their agents were having to solve, but also due to theoretical results on the computational intractability of wide classes of optimal behavior and mechanism design problems (Conitzer and Sandholm 2002)—use of deontic and other modal logics, particularly KD45 (Lomuscio and Sergot 2002) and its variants, to animate agents. Such specifications are rich languages in which to express obligations and norms, things that are expected of individuals and separations between what is permitted and what is not. These logics have very much the flavor of 'behavioral rules' with the realized execution of such specifications being actual agent behavior.

7 Cognition: Sensory Order vs Parallel Distributed Processing

More than once I have been pointed to Hayek's *Sensory Order* as a fertile starting ground for thinking about agent behavior from a bottom up point of view. In that work Hayek seems to have co-discovered Hebbian learning (Hebb 1949), although without the mathematics. He also drew diagrams and made claims for systems that today we would call neural networks, again without any formalism. It is probably too much to say that he anticipated later developments in neural networks insofar as the seminal document in that field *after* Hayek essentially brought the field to a close before it had even started (Minsky and Papert 1969). (Given that today it is conventional to think of neural networks as simply cascaded logistic regression I believe there is much less interest in the subject than once was the case.)

Reading Hayek on this subject from perspective of modern cognitive science we are, once again, seemingly in the position of Hayek having anticipated certain modern developments. But without either significant mathematical prowess or the luxury of the later literature, his work, although prescient, seems to have gone largely unnoticed by the later innovators within psychology and cognitive science. For example, there is no reference to Hayek in the important early cognitive science text of Newell and Simon (1972). I think Hayek's contribution here is less substantive than methodological. Given that this book was written in the heyday of behaviorism in psychology, its focus on brain and mind as bottom up, distributed processes was usefully out-of-step with psychology researchers of his day. It would take a generation for the 'parallel, distributed processing' (PDP) perspective to take off (Rumelhart and McClelland 1986; McClelland and Rumelhart 1986) and even longer for a research program close to Hayek's vision to become popular (e.g., Dennett 1995; Cacioppo, Visser and Pickett 2006).

Today we do have researchers who know and understand what Hayek wrote on this topic, and even some pursuing topics in cognition and consciousness from a perspective commensurate with his. For example, 'The Hayek Machine' is a robot that lives in a 'block world' and which organizes blocks using very little information about its environment. (Baum 1999). More recently, large scale computational efforts to model whole brains 'from the bottom up' (Markram 2006; 2012) is very much in the spirit of Hayek's ideas in this large, broad area, but his actual influence seems to be small.

8 Micro Actions, Macro Patterns: Equilibrium vs Ecology

It is my understanding that there is a long tradition in Austrian economics of resisting equilibrium notions and theorizing (Rosser 2009), although some authors came to a grudging acceptance of equilibrium notions as a pragmatic expedient. Having read Hayek's acceptance speech for the Sveriges Riksbank Prize in Honor of Alfred Nobel it does seem to me that, at least by this late time in his life, he was open to thinking in terms of tendencies of an economy toward equilibrium.

> [I]t is my conviction that if we want to explain economic phenomena at all, we have no means available but to build on the foundations given by the concept of a tendency toward an equilibrium. For it is this concept alone which permits us to explain fundamental phenomena like the determination of prices or incomes, an understanding of which is essential to any explanation of fluctuation of production. If we are to proceed systematically, therefore, we must start with a situation which is already sufficiently explained by the general body of economic theory. And the only situation which satisfies this criterion is the situation in which all available resources are employed.

It seems to me that much of the writing in economics on the distinctions between equilibrium and non-equilibrium (e.g., Kaldor 1972; 1985; Farmer and Geanakoplos 2009; Arthur 2002) is problematical because it fails to distinguish between levels, as in figure 1. As I have written elsewhere (Axtell 2014), there is a sharp dichotomy between agent-level fixed points (micro equilibria) and aggregate stationarity (macro equilibria). Equilibrium at the agent level always implies equilibrium at the macro-level, but the converse is not true. However, agent-level dynamism does not imply aggregate dynamics, for it can be the case that the macroscopic statistics will be in a stationary configuration despite the micro-dynamics.

So imagine a world in which the macro-level is stationary but the micro-level is in a state of perpetual adjustment and dynamism. Agents work hard to make a living and are always trying to outperform one another, sometimes with success and other times not. This is the world of atomic motion but constant temperature, of a billion financial trades each day but little change in price indices, of a dozen engineers moving from Ford to Chrysler while a comparable number go from Chrysler to GM at the same time that number go from GM to Ford. Conventional labor market models postulate this kind of steady-state job flux, but it is hard to write down working models that have it, and still harder to call an economy that features it to be in general equilibrium! Clearly this is an important lacunae of the neoclassical system.

There is a stream of theorizing in complex adaptive systems that asserts that the local level is perpetually out of equilibrium while the aggregate system may achieve something approaching a steady-state (Holland 1998; 1995; 1976; 1993; 2012). This is reminiscent of Hayek but I note that Professor Holland is a computer scientist, with significant training and background in cognitive and behavioral science as well as in biology, who does not cite Hayek anywhere that I can find. So we are left with the impression that in its modern incarnation agent-level disequilibrium seems to owe more to the Holland than Hayek.

In fact, back in the '90s Professor Holland got into a debate with Ken Binmore, the game theorist. For Holland the nature of interactions are deeply non-equilibrium while for Binmore people either play mixed strategy Nash equilibria, correlated equilibria, or some other form of agent-level fixed point. Holland argued against this by analogy to the rainforest, which he described, from an ecological point of view, as in a state of perpetual ebb and flow at the micro level despite being in something like a stationary state at the aggregate level. While Holland's perspective certainly seems preferable to Binmore's, it is also the case that there is little mathematical understanding of these kinds of behavioral ecologies and so they are almost terra incognita at the present time. Yet the metaphor of an economy as an ecology is frequently encountered (e.g., Farmer 2002) and seems, in the main, to be a fertile way of thinking about the distributed, decentralized interactions of large number of individuals.

9 The Veil of Complexity: Gardening vs Ecosystem Management

I argued in the first section above that the Hayekian notion of spontaneous order—arrangements involving people but not explicitly designed by a single person—was a kind of emergent order that also possessed a positive welfare measure. Hayek made it clear that he felt that 'top down' intrusion into well-functioning social institutions was often counter-productive, and so an important postulate for Hayekian public policy was non-interference with the normal operation of existing social orders.

> "If man is not to do more harm than good in his efforts to improve the social order, he will have to learn that in this, as in all other fields were essential complexity of an organized kind prevails, he cannot acquire the full knowledge which would make mastery of the events possible. He will therefore have to use what knowledge he can achieve, not to shape the results as the craftsman shapes his handiwork, but rather to cultivate growth by providing the appropriate environment, in the manner in which the gardener does this for his plants." *The Pretense of Knowledge* (Hayek 1974)

In my book (Epstein and Axtell 1996) we explicitly used the gardening metaphor in our title, *Growing Artificial Societies*, both by virtue of an affinity for the Austrian perspective and because 'growing' really seemed to describe what our models were doing. We started out with a basic environment, our 'soil', modeled as a cellular automata (CA), and to which we added our 'seed' agents, having very simple rules of behavior. We then advanced the model forward in time to see what 'grew up.' Based on that knowledge we then changed the rules somewhat, replanted a population of agents, and ran time forward again. Sometimes, if the results were highly variegated we would have to 'plant many crops' in order to fully understand what was emerging from our computational 'garden'. We could also think ahead about what we wanted to come out but rarely, if ever, could we make that happen without substantial parameter tuning or outright alteration of the rules of behavior. But we were always thinking 'from the top down,' as it were, while building 'from the bottom up,' with complete and precise control over every aspect of our closed world, and still it was hard to know exactly what 'specimens' would manifest themselves.

So the job of an actual policy 'gardener' is much, much harder: she does not know exactly the current state of the garden, has a harder time predicting the effects of policy changes than we did, cannot run experiments repeatedly to see in advance what will happen, cannot perfectly induce her 'agents' to comply with her dictates, and cannot noiselessly measure the result of her policy changes!

So is gardening really the right metaphor? Because of the near impossibility of point predicting the future, the gardener makes no attempt to forecast his garden's exact yield, but rather purposively plants species that are desired and will hopefully produce worthwhile output. There is significant stochasticity in the gardener's yield, but there is usually not literal surprise, i.e., planting snapdragons and getting tomatoes! Perhaps the role of policymakers is closer to forest management, in which there is a complex ecosystem of predators and prey, of failing patches and thriving ones, of easily managed areas and impenetrable, wild ones. While the forest may be in a kind of macro-steady-state the individual trees are each at various stages of their lifecycle, some animals are coming while others are going, with rainfall often undependable. The microscopic structure of the forest is intrinsically dynamic even through at the level of yards of timber available to be produced it is, by virtue of the law of large numbers, stationary. The forest ecosystem *evolves*.

10 Institutions: Origin vs Operation vs Evolution

A gardener grows her garden, is responsible for it, takes credit for good results and is disappointed with bad results. But the word 'growing' does not mean that she creates the garden from whole cloth. She has to start with something, and soil and seeds are the usual points of departure. Distinct sets of seeds produce different kinds of gardens.

When we begin the process of model building, starting with individual agents in some kind of economic environment, we know that certain kinds of rules will produce certain kinds of economies. Rules of exchange beget economic trade. Increasing returns to scale in production induce teams to form. Being able to condition behavior on an observable characteristic may be sufficient to create segregation. Imitating neighbors may produce herding, and so on. Today we know a little about how to get certain kinds of multi-agent institutions—markets, firms, social norms—to grow from the bottom up, we know a little more about how such institutions operate, we know a lot less about how they evolve, about their lifecycles, and how they give way to subsequent generations of institutions, for no single institution, no matter how large and powerful, lives forever.

Institutions are perhaps the most neglected feature of modern economics. While they are paid great lip service, and Nobel prizes are awarded for their study, we do not have, today, anything like a coherent *theory* of institutional creation/formation, growth, and turnover. That is, we know very little formally about the lifecycle of institutions, their internal structures, their governance, the amount of welfare they generate, and so on. Hayek was deeply interested in institutions, and his writing on this subject is typically pregnant with opportunities for building models that would illustrate basic ideas. To see that conventional models neglect institutions, think of the WMAD theory of markets described above. In that theory it is individuals who trade, not firms, although there are various ways to add retail firms into the theory. There are no trading rules when in reality the procedures for acting as a market maker on the NASDAQ run hundreds of pages long. Regulatory and oversight institutions are real but essentially unrepresented theoretically, except perhaps as some crude constraint. Now clearly WMAD is an abstract model, and that fact that it does not include all these real-world complications is viewed as a good thing. But if we systematically do not include institutions at any point in our models how can we say anything about whether or not they make a difference? For example, the burden of bureaucracy, said by many to be rapidly increasing over the years, has no analog whatsoever in any extant field of economics of which I am aware. Widely recognized as mportant but neglected in practice, that is the state-of-the-art when it comes to institutions. Other fields perhaps do better, as in organization science where the in-depth study of single organizations is routine. But such models rarely make their way back into economics.

11 Macro, I: Exogenous Shocks vs Endogenous Dynamics

It is conventional wisdom that the recent financial Crisis was due in large part to the U.S. housing market bubble, which it self was caused by high leverage mortgage lending (Gorton 2010; Commission 2011; Gorton 2012; Geanakoplos *et al.* 2012; Atkinson, Luttrell and Rosenblum 2013). While this assessment is not universally held, most explanations for how the crisis played out describe the unfolding of events in terms of components of the overall financial system, i.e., banks, hedge funds, corporations and their credit needs, households and their credit card debt, etc. What I am arguing here is that there is a heavy-duty *endogenous* component to the Crisis that essentially involved the machinery of the financial systems turning on itself and amplifying whatever the initiating event was into a full-blown panic. While *exogenous* events may have further amplified the situation, much of the damage was caused by the *internal* linkages and institutional structures that characterize the financial system today.

The concept of endogenous dynamics is not one that is easily swallowed by mainstream macroeconomists today, builders and users of DSGE models that postulate the microeconomy to be in a state of perpetual equilibrium, with all perturbations of such equilibria due to exogenous, stochastic forcing. Earlier generations of macro models, including Austrian capital theory, dealt more easily with out-of-equilibrium dynamics. So it is unclear to me whether increased mathematical sophistication in macro has yielded real progress.

As argued above, it seems to me that a complexity economics worthy of the name replaces agent-level equilibria with dynamics, and then the macro level is either in a steady-state or not, depending on the nature of the overall economy. The exact reasons why sometimes fluctuations are amplified and sometimes not is today opaque, partly because it is outside the scope of those primarily concerned with exogenous economic drivers. In fact, the real economy and the financial sector are each so large and variegated that it seems extraordinary we do not experience more dynamism than we do. That is, with financial markets largely unencumbered by regulations it is a bit surprising they have as much stability as they do. Indeed, a generation ago Robert May studied the stability of ecosystems (May 1974) as they got larger and came up with the surprising hypothesis that perhaps largesse was associated with *in*-stability.

While I am not aware of general results on this topic, two recent Mason Ph.D. theses in Economics illustrate the difficulties of finding equilibrium points as the scale of an economy increases. Nathan Smith wrote a dissertation on competition theory, building a model in which he could let the user specify the number of firms, n, moving from oligopoly ($n = 2$) to pure competition (n large). Painting with a coarse brush—his dissertation runs several hundred pages—he found that for $n \sim O(10)$ the dynamics of supply and demand were so complicated that he had to program reasonably sophisticated statistician agents into the behavior of each firm, in order to give them any chance whatsoever of discerning actual supply-demand conditions and making production plans accordingly. But even with complex firms the economy would display fairly elaborate dynamics and never really settle down to anything like a fixed point equilibrium.

In a different domain but with comparable qualitative results, Will McBride built a model of private money issue by banks, roughly comporting with the Scottish experience in the late 18th Century. He found that with just a couple of currencies he could get stable, reasonable results, but that as soon as he increased even a little beyond this level, to 5 bank-issued currencies, say, that the dynamics quickly became very complex. In fact, given that his primary hypothesis was that non-State money would be superior than normal State-issued currencies, the intrinsic complexity of his results largely led to the rejection of the original hypothesis.

12 Macro, II: Mathematical Economics vs Mathematical Tractors

It seems to me there are at least three possibilities for the future of macroeconomics. First, microeconomic activities aggregate cleanly into coherent macroeconomic wholes, which have a definite relation to one another, independent of the microeconomics. This would represent a kind of 'statistical mechanics' for the economy: in reality the air molecules buzzing around the room at the speed of sound combine to yield the overall air temperature, volume and density in a way-well-described by the ideal gas law. We can then talk about how to make the room warmer or cooler based solely on this aggregate relation without reference to the air molecules and their properties. Of course, there is an explanation of what happens at the macro-level in terms of the micro-dynamics, it is just that we do not need to go all the way down to that level to understand how to make the room warmer or colder.

Secondly, it might be the case that there do not exist very useful aggregates, sufficiently well-defined, stable and meaningful that a closed form mathematical theory inter-relating only the aggregates exists. If this is true then the quest for microfoundations of macro is not merely desiderata but a necessity, for the only way to properly understand such an economy would be to start from the current macro-state, disaggregate down to the current micro-state, march the microeconomy forward in time, and reaggregate. This is the economic analog of Coleman's boat (1964). If this turns out to be the best we can do with macroeconomics then I suppose a fair question to ask is whether a true macroeconomics is possible? That is, macroeconomics with microfoundations of the 'U' variety is really just microeconomics, with occasional aggregation, plain, pure and simple.

In order to realize this second possibility it has to be the case that, once we are at the micro-level, we can 'spin' the economy forward with a model of some kind. If we cannot do that then the third possibility is that macroeconomics, whether written in terms of aggregates or individual relations, does not exist. It seems to me that most economists believe the second of these possibilities, while minorities are in the first and third camps, an armchair empirical claim. Perhaps someday economics as a science will progress to the point that it is clear which of these possibilities is true. Here I want to address a slightly different point than which of these three possibilities is likely to be scientifically correct.

Consider another elaborate 'machine', say a tractor. It is complicated because it is made of a wide variety of interacting subsystems—fuel, directional control (steering), weight/balance (frame), power creation (engine), power consumption (useful work, e.g., plowing), electrical (ignition and gauges), speed controls (throttle and brakes), power train (transmission), operator comfort (seat, cabin) and so on. The interactions between these distinct systems can lead to a tractor operating at high performance or poorly. But although a tractor is complicated it is not complex in the sense of being able to produce novelty, emergence, and so on. It is designed from the top down by eminently practical people (engineers) for other practical people (farmers). So why is it that in the owner's manual or operating manual for a new tractor there are no equations representing how it works, where its stability points are, how its performance can be optimized? Now one answer could be that such abstract representations of the tractor's 'behavior' are not given because there are few farmers who would make use of it. But I do not think that is why no such representation is made available to users or even published in scientific journals. The real answer, I believe, is that even within the bowels of John Deere Engineering there does not exist a closed-form, analytical model of a tractor. It does not exist because to put such a model together would be a *very* difficult undertaking, with some reward certainly, but the costs probably outweighing the benefits. Now arguably, a tractor does not need a mathematical model in order to be useful because it has been designed with a different user interface, its cab. Here the day-to-day, minute-to-minute, even second-to-second operations of the tractor are handled by a single mind. This interface has been designed to give the user a great deal of both control and flexibility in operating the tractor, so there is no real need for any equations.

An economy is certainly more complex than a tractor! As a practical matter a typical economy includes within it millions of tractors, so the economy cannot be less complex that the tractors which are part of it! But if there is no practical/useful way to write down equations for a tractor then why do we think it is a useful thing to do for an economy as a whole? Writing down equations may not be useful to the people who 'operate' the economy, whether these are chief executives, Federal Reserve Board Governors, or an investor from Main Street. In particular, it may be that having a good user interface for the data coming in about the economy—a dashboard—and control panels for modifying rules and regulations that affect the economy—like a NASA mission room—would go a long way toward satisfying most of the needs of policy-makers. Today we know a lot about how to 'eat' large data streams and condense them into intelligible forms. A future macroeconomics that recognizes the primacy of the data and is humble about the ability to predict may look a lot more like a rocket launch room than the wood-paneled offices and cloistered meeting rooms of a typical central bank.

In the late 1920s a natural scientist named at Oxford named Frederick Soddy, with no reason to be humble because he had won the Nobel Price in Chemistry for his work on radioisotopes, recognized that the economics of his day was data poor, and that the policies being proposed where thus potentially wrong. He looked at all the data he could find and concluded that the way out of the difficulties was for Britain to:

i. Form national economic data gathering institutions;
ii. Let exchange rates float;
iii. Abandon the gold standard;
iv. Use budget surpluses and deficits as counter-cyclical policy tools;
v. Stop banks from creating money and debt.

These were each extremely radical proposals for the time and he was immediately labeled a 'crank' for proposing such heresies. The better part of a decade later Keynes would pay him faint praise as one of "…the brave army of heretics…who, following their intuitions, have preferred to see the truth obscurely and imperfectly than to maintain error, reached indeed with clearness and consistency and by easy logic, but on hypotheses inappropriate to the facts" (Keynes 1936).

13 M-foundations of (M+1)-economics: Micro, Meso, Macro

For the past generation the quest for satisfactory microfoundations of macro has been a kind of crusade for the holy grail. Here I suggest, based on the discussion of institutions above, that the only reason conventional macroeconomists look for microfoundations is because neoclassical economics completely neglects all intermediate levels. At the very least macroeconomics should seek mesoeconomic foundations, and then mesoeconomics needs microfoundations. And at each one of these levels we really should be talking about emergent economic structures.

As a concrete example, consider a 3-level economy in which agents form firms and firms constitute the main part of the economy. So here the meso level is the firm level. In this U.S. this means 6 million businesses with employees, 120 million workers total, and another 25 million farms, mom-and-pop stores and other operations having no employees (Axtell 2013). This alternative type of macro model would certainly not be 'solved' by rational expectations since (a) the data on firms makes very clear that typical businesses experience regular and large fluctuations from year to year, certainly beyond anything that might be explained by rational deductions, and (b) a solution at the macro-level would not involve the decision making of individual workers two levels down—that is at the wrong level in a 3-level economy; elementary particles (e.g., quarks) do not show up in physical chemistry models and worker effort level decisions would not show up in a macro-model with firms at the meso level.

If reasoning across two levels can be problematical, as has been argued at the start of this essay, the dual fallacies of division and composition really manifest themselves when reasoning across many levels. Specifically, while it may seem reasonable to outfit the representative agent with elaborate rational expectations-type calculations to perform over riskless and risky instruments while earning wages and saving for retirement, in a many-level model with institutional layers many of these calculations are things the institutions do for the agents: mutual funds execute investment decisions, wages are determined at the firm level, and so on. All economists know this but DSGE models do not reflect this.

Given the analytical difficulties of building and solving many-level economic models, like the difficulty of putting together a single model embedding all the sub-systems of a tractor, it may turn out that the only viable way to make progress is computationally. If so, agent computing has the great advantage that it scales up linearly, takes advantage of all compute resources (e.g., CPU, memory, harddisk) so that the more you have the bigger and faster your middle levels can be, and now today, for the first time, we can really render models with $O(10^8)$ agents. Among the many competing research programs in macro underway today, all attempting to rejuvenate a subfield that has obviously floundered over the past decade, large-scale agent computing approaches are moving forward (Axtell 2006).

14 Scientism: Econophysics vs Neoclassicism vs Austrianism

I have to admit that my favorite book within the Hayekian oeuvre is a thin volume that has few pretensions toward positive economic science, but rather is a polemic on methodology. I refer to *The Counter-Revolution of Science* (Hayek 1955), the epithetic subtitle of which—'Studies in the Abuse of Reason'—leaves little doubt that its author is not going to pull any punches.[2] In this book Hayek assaults various thinkers, mostly from the 19th C, for their crude attempts to manage or plan their way to more perfect societies—so-called 'scientific socialists' were certainly among his targets—as they had little experience managing real social systems and, most importantly, according to Hayek, were not used to thinking about societies as supple, organic entities. Usually these were people with scientific training, possibly sociologists (e.g., Comte), whose 'habits of mind' were quite different, Hayek asserts, from what is needed in the social sciences. These people tended to see the world as nearly infinitely adaptable to their wishes and predilections so that establishing new laws, legislation and performance rules would not carry with them the possibility of unintended side effects and consequences. Such people are too sure of themselves,

[2] In this it is not unlike Veblen's remarkable *The Higher Learning in America: A Memorandum On the Conduct of Universities By Business Men* (1918) which, according to Heilbroner (1980), was originally to be subtitled 'A Study in Total Depravity'. In comparison to Veblen and Hayek, academic authors today seem to have a lot less liveliness in their titles, a deficit in panache.

without sufficient appreciation for the complexity of social structures, Hayek asserts, and their interventions rarely work out as desired. For Hayek, the counter-revolution of science is that people with scientific credentials end up mucking up, at best, or downright destroying, at worst, the normal form and function of self-organized societies. Early socialists like St. Simon and Fourier, quantitatively oriented social scientists like Pareto and Gini (for whom many Austrians had little appetite, considering them the high priests of fascism (Gini 1927)), and especially students of L'Ecole Polytechnique, all of these people are exposed to Hayek's withering critique.

I like the temperament of this book because it is so unabashed. And during the 20th C it is certainly the case that followers of these schools of thought created much havoc. But it does seem to me that it is not the scientists among them who did the damage, but rather the politicians and policy-makers who adopted such hair-brained schemes, whether Lenin, Mussolini, or Mao.

I sometimes think that there are two kinds of economists, dual to one another. One kind wants to minimize costs, wring out all inefficiencies, doing more with less, make the public sector work like the private, and so on, often a humble, earnest, indefatigable, even zealous person. The other cares more about increasing profits, being innovative, trying new things, and cuts a wide path through life. Each can contribute substantively the economy but they have almost nothing in common.

Similarly, it seems to me that there are two distinct kinds of 'scientistic' attitudes and they are not equally objectionable. When natural scientists stumble, for whatever reason, into economics and finance, they often bring with them substantial data analysis experience, some computing skills, and perhaps experience working for a financial firm. These abilities can contribute to solving important economic problems, especially to the extent that large datasets need to be analyzed and computer models built, as is common in finance. For example, for the last 20 years or so physicists have been wandering around finance, trying to bring concepts and tools from condensed matter and other complexity-related fields within physics to bear on financial data (Mantegna and Stanley 2000). Some working within this emerging 'econophysics' tradition have even gone so far as to declare economics to be a new-fangled physical science (Farmer, Shubik and Smith 2005). I think it would be very easy to hurl Hayek's 'scientism' charge at these people but I feel relatively strongly that it would not be helpful to do so. For these are scientists who have no particular political axe to grind and are not captives of some political party with wild ambitions. Rather, they are people with deep scientific experience who see the way certain economists do science as shallow, data poor, and not the way they would proceed. So they try to do it differently, with some significant successes (Lux 1998; Lux and Marchesi 1999; Farmer 2002; Farmer, Patelli and Zovko 2005; Cont, Ghoulmie and Nadal 2005; Cont 2006). The econophysicts' critique of mainstream finance, with its utility maximizing, highly rational traders, is that such models simply do not comport with their experience at trading desks, writing code for high frequency, algorithmic trading, off the floor trading in 'upstairs' markets, and so on. So they build simpler models, near zero-intelligence (ZI) trader models that do as well or better than neoclassical market microstructure models empirically. They try to fit data. Overall, their contribution is surely positive. But they typically do not model behavior very deeply. They are good on taking the complexity of real markets into account but use the institutional structure of markets to constrain their models, not complex agents. Neoclassical finance gives short shrift to the real complexity of markets while attempting to represent behavior mathematically. And against each of these two approaches to finance and economics, Austrians recognize the complexity of the real world and describe economic processes with behavioral richness, but do not build models. I summarize these observations in figure 10.

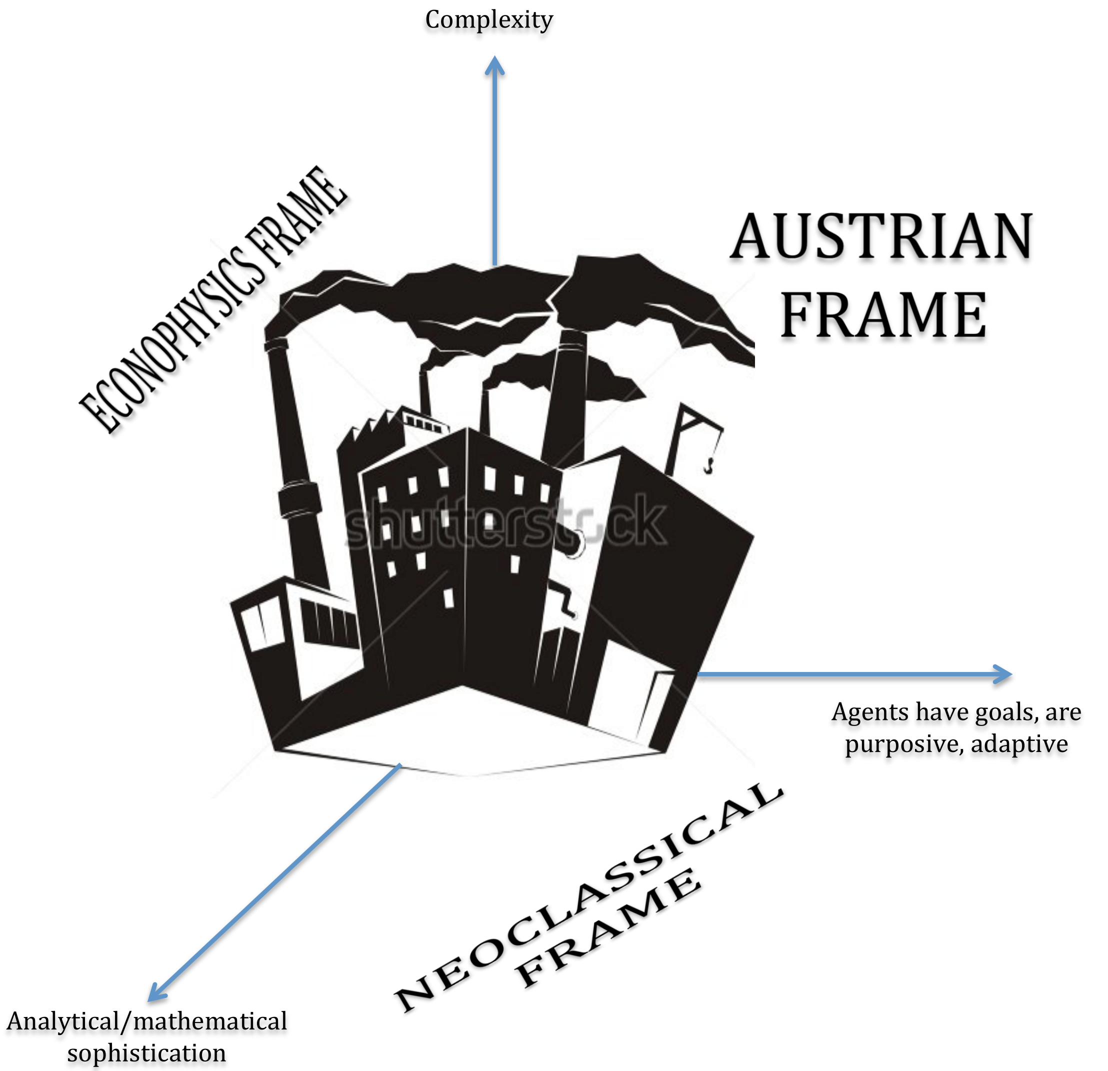

Figure 10: The real-world economy lives in a multi-dimensional space associated with high levels of complexity, agents who are inherently purposive and adaptive, and which generates quantitative data suitable for calibrating models. Neoclassical economics replaces real-world dynamism and heterogeneity with equilibrium postulates and representative agent models. Econophysics replaces agent-level behavior and goal-seeking by stochastic processes. Austrian economics has, historically, been little willing to enter models into the scientific arena.

Now, these 'frames' are obviously caricatures, as there exists some neoclassical work that tries to take real-world complexity into account, there is some research in the Austrian tradition grounded in models, and certainly there are some econophysics papers that explicitly represent agent behavior, as my neoclassical collaborators, Austrian colleagues, and econophysics friends will surely tell me. But in the main each of these framings emphasizes some aspects of the fully complex economy while leaving others out.

II Prospective: Coevolution of Complexity and Austrian Econ?

"*A long-standing tradition makes the first step towards introducing order into the plethora of the phenomena encountered by man by classifying them as either 'natural' or 'artificial.' Obvious and sensible as this classification may appear to be, it is deficient. Characterizing 'natural' phenomena as those that are wholly independent of the human sphere and 'artificial' phenomena as those that involve human agency leaves the deficiency at least partially concealed. It is brought into relief, however, once it is said that the realm of the 'natural' consists of everything that is the result neither of human action nor of human design, while the real of the 'artificial' of everything that is the result of both: what is clearly seen to be lacking is the at least hypothetical realm of everything that is the result of human action but not of human design*."

Edna Ullmann-Margalit (1924-2007)

Conventional, neoclassical economics is essentially mathematical in its methodology, severely constraining the realism of its models. Complexity economics utilizes instead agent-based computational methods to represent boundedly rational individuals having local information acting out of equilibrium. The more supple expressiveness of agent computing means that higher fidelity models can be built, including ones that capture more of reality that the simple neoclassical ones. For much of its history Austrian economics has recognized the limitations of the mathematical approach and has eschewed formal models. While this has limited the penetration of Austrian ideas into the mainstream of the profession, the new tools and techniques emerging from the complex systems literature would appear to offer Austrians new opportunities to illustrate their ideas with models, to animate their market process descriptions with trading agents, to leverage the 'manna from heaven' that is exponentially-growing computing power to better understand the economy.

What I have tried to argue above is that not only does Hayek deserve credit for beating the physicists and other complex systems scientists to the punch in articulating a coherent view of complex systems, his radically distributed and decentralized view of the world is a wellspring for renewal of the complexity program as the methodology continues is colonization of new scientific fields. As a practical matter it is my experience that, for people who are new to complex systems, some of the most common mistakes made are not taking a consistent bottom up perspective, worrying too much about comporting with neoclassical norms (e.g., market clearing), and not letting emergence to do its work, letting order (and disorder) arise naturally—newcomers to the field want to 'build in' their results, especially when arriving with a lot of neoclassical 'baggage'. Building good agent models means taking the inherently distributed, decentralized, asynchronous social world at face value and avoiding doing 'unnatural' things that agents in the real world never do, like sorting an entire population, participating in Gale-Shapley matching, tessellating a price simplex, and other computationally intensive tasks that have no analog in the real world. Against these bad behaviors Hayek is an antidote. Among the Hayekian lessons worth learning:

1. Spontaneous order is emergence with a welfare function.
2. Tacit value can reside both intra-agent (knowledge) as well as inter-agent (relationships).
3. The conventional theory of general equilibrium is sufficiently incredible computationally that almost any market process story is preferable.
4. 'More can be different' and sometimes the idiosyncrasies of individuality do not matter, i.e., certain kinds of processes have the 'universality' property.
5. Individuals are neither atoms nor robots and most are not sociopaths.
6. Behavioral rules are an alternative to $max\{U\}$ specifications and evolve.
7. Perception and cognition can be thought of as bottom up processes.
8. Agent level dynamism implies ecologies but possible macro level stationarity.
9. We cannot know the exact current or future state of a garden, much less an economy, so management of either requires an evolutionary approach.
10. Few institutions play important roles in models today.
11. Macroeconomic dynamics are endogenous to the economy.
12. Equation-based macroeconomic models may be too simplistic to be useful.
13. Macroeconomics needs mesofoundations and mesoeconomics needs microfoundations.
14. Scientism is less a risk than an opportunity today.

The upshot of all this is that modern day Austrians can adopt the emerging tools of the complexity economists while the latter can profit from adopting a purer Hayekian perspective. Can the two distinct communities work together, coevolve? Can there be 'team production' so that the whole is more than the sum of the parts? What are the barriers to collective action? How does one 'ignite' a virtuous circle, a positive feedback loop? Can this be done at one institution (e.g., George Mason), or does it require a larger community? What is to be done?

III Proscriptive: The Road to Academic Serfdom in the Age of Models

"[W]hen you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind; it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science, whatever that may be." William Thomson, Lord Kelvin (1824-1907)

Science is a difficult business. It is competitive. It forces its performers to adopt unnatural, even wild views, ones that may run very much against their personal predilections (think quantum mechanics for counter-intuitive, or your least favorite behavioral 'anomaly' for assaulting preconceptions). Science evolves, but individuals are not set up to evolve, that is what populations do—thus Planck's maxim that "Science advances one funeral at a time," and Laughlin and Pines. So if a scientist is unwilling to modify his or her understanding of how the world works, their career will be short-lived. Mostly, science is hard because Nature does not readily reveal her secrets, her modes of operation, her crown jewels as it were.

Indeed, it was only 100 years ago that Frederick Soddy first figured out how radioactive decay transforms one element into another, the impossible Bohr atom ruled the roost in physics, biologists did not understand what chromosomes were made of, and commercial radio did not exist. How far science has progressed!

Which brings us to economics. I am well aware that not all economists believe our subject is a science. I leave it to the philosophers of science to adjudicate which side has the stronger arguments today. But perhaps we can all agree to entertain the hypothesis that economists can *strive* to be scientific, and that someday the profession might develop to such an extent that it could reasonably be considered a science. Whether one thinks of economics as a science or not, I believe there is widespread agreement that, like the natural sciences, doing research in economics is not easy. Economic data are often very noisy, experiments are a relatively new methodology and typically limited to small scale, and mainstream theory is either simplistic as a description of individuals with 100 billion neurons or else assumes people are omniscient, unconstrained by their cognitive apparatus.

Whether one thinks economics is a science or not, it is surely a field whose dominant methodology today is the systematic use of mathematical models. One need only crack open a recent issue of a leading journal to find papers full of theorems, corollaries, propositions, and their associated proofs, tables of regression results, numerical solutions, mathematical programs and/or simulations. It is not a field where much progress can be made with verbal arguments alone. Rather, economics is a *quantitative* discipline in which *qualitative*, verbal argumentation is rarely decisive. While insights and useful pedagogy may result from exploring the history of economic thought, performing exegesis on classical texts, or resurrecting defunct policy prescriptions from long dead authors, today in the economics profession *it takes a model to beat a model*. It is tempting to say 'it takes a theory to beat a theory' but an extant theory, embodied in a model, cannot be beat with a theory for which no model exists. I do not know if this is a satisfactory state of affairs but it is how the world of economics research works today.[3]

If Hayek were alive today would he be happy with the current position of Austrian economics? Would he see the progress that has been had? Would Hayek be an Austrian? (I am reminded that even in Marx's lifetime he disavowed Marxism.) I believe that in order for Austrian economics to become more relevant in the economics profession it has to become more cumulative, with the findings of one research team being used by another to progressively build up a core group of economic…models. Let's formulate an alternative to the WMAD model that is consistent with market process ideas and study its properties, highlight its strengths, and get behind it as a new standard. Let's create canonical models that can be used to study the emergence of money, its role in a modern economy, and its evolution. Let's create models that generate spontaneous orders of the type Austrians care about, for purposes of studying their robustness, resilience, and overall performance. Under what conditions is self-governance stable? What kinds of

[3] It is my impression, from occasional interactions with philosophers of science, that they view the elimination of 'history of thought' in a discipline as an essential signal that a field has become scientific, thus bearing to some extent on my previous remarks on the scientific status of economics.

endogenous behaviors or exogenous shocks can destabilize spontaneous orders? Let's 'grow' a whole (compute) farm full of spontaneous orders and study their properties from a population perspective: what are their lifetimes, can they evolve too fast, what makes some work better than others? I am not aware of any substantive body of work on any of these topics although doubtless there is some!?! Could it be that up until now we have been living in a kind Kelvin-class prescientific era of Austrian economics. The time is ripe to leverage the new tools. The bleak alternative is self-segregation to the periphery.

IV Prescriptive: It Takes a Model to Beat a Model

Neoclassical economics is a simplistic view of a complex subject. It represents a kind of 0th–order approximation to reality. Doubtless, some of its tenets are true, but 1st and 2nd order corrections are likely to strongly modify or even outright reverse the policies that spring from the basic models. While formal methodologies within economics come and go—mathematical programming giving way to optimal control theory replaced by game theory and now experiments—and social norms of exposition evolve—I believe there are more theorems, corollaries, propositions and proofs appearing in the *AER* now than at any time in its history, on a per article basis—Austrians have participated little in these developments as Austrian methodology is little changed from the time of Hayek, being largely philosophical in tone and character.

The Austrian 'projection' into model-free space can be viewed as a reaction to the crude methodologies model-building economists have wielded to date, just as the econophysicists' projection into behavior-free space comports with their unwillingness to commit to the standard neoclassical behavioral (max{U}) model. Neither Austrians nor econophysicists are willing to accept the Faustian bargain of jettisoning economic complexity, as the neoclassicals have done and who are now forced to live in an unrealistically simple economic world.

But things are changing, new tools and techniques are arriving, a fresh wind is blowing that will carry science to new worlds. For we live in the era of *computational exploration*. This happens once, like European exploration of the Earth by sail. Small sailboats, say corvettes, cannot make it around Cape Horn, sailing from Plymouth or Porto. And once all the world has been 'touched' there is less 'quest' in further exploration. The first few trips to Easter Island are noted in the history books, the initial survey of the Galápagos Islands revolutionizes biology. Today enough computing power can be marshaled to simulate (1) how the Moon formed (Cuk and Stewart 2012; Canup 2012), (2) the incredible (and beautiful) motion of whole spiral galaxies with every star represented {{Binney, 2008 #2505}}, (3) the medical function of individual molecules in advance of laboratory synthesis (Lewars 2011), (4) the entire chemical function of a cell (Karr *et al.* 2012), (5) every neuron in a brain (Markram 2012), (6) whole ecosystems (Railsback and Grimm 2012), (7) epidemic control strategies at national scale (Gemann *et al.* 2006), (8) simple anthropologically-significant societies at full-scale over thousands of years (Axtell *et al.* 2002), (9) the flow of workers between jobs at country scale, i.e., every firm, every worker (Guerrero and Axtell 2013), and (10) the reshaping of political borders by military power (Cederman 2002), to name just a few prominent examples. Elsewhere I have wondered what the social sciences would be like today if Smith or Malthus or Marx had had a modern computer. Perhaps Austrian economics would look a lot different today if the founding fathers had been able to truly experiment with how individuals create order spontaneously or by building agent models of markets as they observed them. Perhaps we would be a lot farther along in the enterprise of understanding, scientifically, how economies work if today we had 50 or 100 years of experience building agent-based models. Somewhat unfortunately, we are all relative newbies to these new tools and techniques. But the good news is that most of the good work has yet to be done!

This era will be known in intellectual history for the birth of large-scale agent-based computing. Today we have lots of speculative, preliminary models but few that seem canonical. We know a fair amount about how to model markets with agents, less about firms, almost nothing about institutions, and essentially zero about the economy as a whole. These are all opportunities. Careers will be made building models on the origin of the family, private property and the state. Exciting times are ahead!

I conclude, unsurprisingly to those who have followed this far, that Hayek was *ahead* of his time in his thinking about the economy as a complex system. The lack of formalism in his work, which some may consider to be a 'bug' is actually a 'feature' from the perspective of building new models based on his insights. I do wonder which researchers are in the best position to accomplish this, complexity-oriented modelers who read Hayek or Austrian scholars who can program. While there has probably been some self-selection in the Austrian ranks, with people not disposed to mathematical exposition choosing to work within the Austrian traditions simply because it was possible to do so, the exponential growth of computing power promises to progressively lower the 'entry costs' to coding and should make agent-based computing more routine and less of an art in the future. After all, we are economists and as the price of a good falls we should consume more of it, right? And in the long run if Austrian ideas are developed into cumulative models by people with quantitative skills, the self selection pressures within the Austrian school would themselves evolve.

I am very confident that Austrian economics of the future will be rich with models, for otherwise there will probably be no Austrian economics. The economics profession needs researchers who are not wedded to its current maximizing ontology. It needs people who are skeptical of mathematical 'rigor' that neglects realism. It needs people who are unwilling to accept extant ideas merely because they are the norm, conventional, the way things have always been done. The field needs skeptical minds, something the Austrian community has in spades. But true progress requires stating ideas in terms of models, and increasingly models are written in code.

The physicist Feynman (1985) once wrote that it is a pity that most of humanity has little or no mathematics. It limits the depth of understanding that non-physicists can have about the physical world. For those of us who studied mathematical economics, essentially taking a math course every semester in undergraduate school (e.g., differential equations, real analysis, topology) and then again in graduate school (e.g., functional analysis, stochastic processes, graph theory), it is tempting to say the same thing about the economic world, that with deep mathematical preparation the way the economy works can not be satisfactorily described. But this is not right, because economies are composed of large numbers of distributed, decentralized, interacting processes of production, exchange, communication and coordination, which evolve over time and are themselves created and operated by purposive, boundedly rational individuals. Today there is no mathematics capable of describing such an economy. Tools and techniques from computer science come a lot closer. Hayek scholars can progressively add to the research program he advanced by adopting these methods.

References

Aaron, H. J. (1994). "Public Policy, Values, and Consciousness." Journal of Economic Perspectives **8**(2): 3-21.

Anderson, P. W. (1972). "More is Different." Science **177**(4047): 393-396.

Anderson, P. W., K. J. Arrow and D. Pines, Eds. (1988). The Economy as an Evolving Complex System. Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume V.

Ando, A., F. M. Fischer and H. A. Simon (1963). Essays on the Structure of Social Science Models. Cambridge, Mass., MIT Press.

Arrow, K. J. and G. Debreu (1954). "Existence of an Equilibrium for a Competitive Economy." Econometrica **22**(3): 265-290.

Arthur, W. B. (1991). "Designing Economic Agents That Act Like Human Agents: A Behavioral Approach to Bounded Rationality." American Economic Review **81**(2): 353-359.

--- (2002). Out-of-equilibrium economics and agent-based modeling. Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics. K. Judd and L. Tesfatsion. New York, N.Y., North-Holland.

Atkinson, T., D. Luttrell and H. Rosenblum (2013). How Bad Was It? The Costs and Consequences of the 2007-09 Financial Crisis. Staff Papers. Dallas, Texas, Dallas Fed.

Axtell, R. L. (2001). "Zipf Distribution of U.S. Firm Sizes." Science **293**(5536): 1818-1820.

--- (2005). "The Complexity of Exchange." Economic Journal **115**: F193-210.

--- (2006). Multi-Agent Systems Macro: A Prospectus. Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model. D. C. Colander. New York, N.Y., Cambridge University Press.

--- (2013). Endogeneous Dynamics of Firms and Labor with a Large Number of Simple Agents, George Mason University: 50.

Axtell, R. L., J. M. Epstein, J. S. Dean, G. J. Gumerman, A. C. Swedlund, J. Harburger, S. Chakravarty, R. Hammond, J. Parker and M. T. Parker (2002). "Population Growth and Collapse in a Multiagent Model of the Kayenta Anasazi in Long House Valley." Proc Natl Acad Sci U S A **99**(supplement 3): 7275-7279.

Axtell, R. L., J. M. Epstein and H. P. Young (2001). The Emergence of Classes in a Multi-Agent Bargaining Model. Social Dynamics. S. N. Durlauf and H. P. Young. Cambridge, Mass./Washington, D.C., MIT Press/Brookings Institution Press: 191-211.

Ball, P. (2004). Critical Mass: How One Thing Leads to Another. New York, Farrar, Strauss and Giroux.

Barbieri, F. (2013). "Complexity and the Austrians." Filosofia de la Economia **1**(1): 47-69.

Bardeen, J., L. N. Cooper and J. R. Schrieffer (1957). "Theory of superconductivity." Physical Review Letters **108**(5): 1175-1204.
Baum, E. B. (1999). "Toward a Model of Intelligence as an Economy of Agents." Machine Learning **35**(2): 15-185.
Beinhocker, E. (2005). The Origin of Wealth: How Evolution Creates Novelty, Knowledge, and Growth in the Economy. Cambridge, Mass., Harvard Business School Press.
Bell, A. M. and W. A. Sethares (2001). "Avoiding global congestion using decentralized adaptive agents." IEEE Transactions on Signal Processing **49**(11): 2873-2879.
Bell, A. M., W. A. Sethares and J. A. Bucklew (2003). "Coordination failure and congestion in information networks." IEEE Transactions on Signal Processing **51**(3): 875-885.
Bird, R. B., W. E. Stewart and E. N. Lightfoot (1960). Transport Phenomena. New York, N.Y., John Wiley & Sons.
Bonabeau, E., M. Dorigo and G. Theraulaz (1999). Swarm Intelligence: From Natural to Artificial Systems. New York, N.Y., Oxford University Press.
Bowles, S. (1998). "Endogenous Preferendes: The Cultural Consequences of Markets and other Economic Institutions." Journal of Economic Literature **XXXVI**(March): 75-111.
Buchanan, M. (2007). The Social Atom: Why the Rich Get Richer, Cheaters Get Caught, and Your Neighbor Usually Looks Like You. New York, N.Y., Bloomsbury.
Cacioppo, J. T., P. S. Visser and C. L. Pickett, Eds. (2006). Social Neuroscience: People Thinking about Thinking People. Social Neuroscience Series. Cambridge, Mass., MIT Press.
Caldwell, B. J. (2004). Hayek's Challenge. Chicago, Ill., University of Chicago Press.
Camerer, C. (2003). Behavioral Game Theory. Princeton, N.J., Princeton University Press.
Camerer, C., T. H. Ho and J. K. Chong (2002). "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games." Journal of Economic Theory **104**(1): 137-188.
--- (2004). "A cognitive hierarchy model of games." The Quarterly Journal of Economics **119**(3): 861-898.
Canup, R. M. (2012). "Forming a Moon with an Earth-like Composition via a Giant Impact." Science **338**(6110): 1052-1055.
Casti, J. L. (1994). Complexification: Explaining a Paradoxical World through the Science of Surprise. New York, N.Y., HarperCollins Publishers.
Cederman, L.-E. (2002). "Endogenizing Geopolitical Boundaries with Agent-Based Modeling." Proc Natl Acad Sci U S A **99**((suppl. 3)): 7296-7303.
Coleman, J. S. (1964). Introduction to Mathematical Sociology. Glencoe, Ill., Free Press.
Commission, F. C. I. (2011). The Financial Crisis Inquiry Report, U.S. Government.
Conitzer, V. and T. Sandholm (2002). Complexity of Mechanism Design. Proceedings of the Uncertainty in Artificial Intelligence Conference. Edmonton, Canada.
Cont, R. (2006). Volatility clustering in financial markets: Empirical facts and agent-based models. Long Memory in Economics. A. P. Kirman and G. Teyssiere. New York, N.Y., Springer.
Cont, R., F. Ghoulmie and J.-P. Nadal (2005). "Heterogeneity and feedback in an agent-based market model." Journal of Physics: Condensed Matter **17**(14): S1259-S1268.
Cuk, M. and S. T. Stewart (2012). "Making the Moon from a Fast-Spinning Earth: A Giant Impact Followed by Resonant Despinning." Science **338**(6110): 1047-1052.
Debreu, G. (1959). Theory of Value - An Axiomatic Analysis of Economic Equilibrium. New Haven, Conn., Yale University Press.
Dennett, D. C. (1995). Darwin's Dangerous Idea: Evolution and the Meanings of Life. New York, N.Y., Simon & Schuster.

Doolen, G. D., Ed. (1990). Lattice Gas Methods for Partial Differential Equations. Santa Fe Institute Studies in the Sciences of Complexity, Volume IV. Reading, Mass., Addison-Wesley Publishing.

Durlauf, S. N. (2012). "Complexity, economics, and public policy." Politics, Philosophy & Economics **11**(1): 45-75.

Eigen, M. and P. Schuster (1979). The Hypercycle: A Principle of Natural Self-Organization. New York, N.Y., Springer-Verlag.

Epstein, J. M. and R. Axtell (1996). Growing Artificial Societies : Social Science from the Bottom Up. Washington, D.C./Cambridge, Mass., Brookings Institution Press/MIT Press.

Epstein, S. (1983). "Aggregation and Beyond: Some Basic Issues on the Prediction of Behavior." Journal of Personality **51**: 360-391.

Farmer, J. D. (2002). "Market force, ecology and evolution." Industrial and Corporate Change **11**(5): 895-953.

Farmer, J. D. and J. Geanakoplos (2009). "The Virtues and Vices of Equilibrium and Complexity." Complexity **14**(3): 11-38.

Farmer, J. D., P. Patelli and I. I. Zovko (2005). "The Predictive Power of Zero Intelligence in Financial Markets." Proc Natl Acad Sci U S A **102**(6): 2254-2259.

Farmer, J. D., M. Shubik and E. Smith (2005). "Is Economics the Next Physical Science?" Physics Today(September): 37-42.

Feigenbaum, M. J. (1978). "Quantitative Universality for a Class of Nonlinear Transformations." Journal of Statistical Physics **19**: 25.

Fermi, E., J. Pasta and S. Ulam (1955). Studies of Nonlinear Problems. Los Alamos, N.M., Los Alamos National Laboratory.

Feynman, R. (1985). Surely You're Joking, Mr. Feynman. New York, N.Y., W.W. Norton & Company.

Frisch, U. (1995). Turbulence: The Legacy of A.N. Kolmogorov. New York, N.Y., Cambridge University Press.

Fuqua, W. C., S. C. Winans and E. P. Greenberg (1994). "Quorum sensing in bacteria: the LuxR-LuxI family of cell density-responsive transcriptional regulators." Journal of Bacteriology **176**(2): 269.

Gabaix, X. (2011). "The Granular Origins of Aggregate Fluctuations." Econometrica **79**(3): 733-772.

Gard, T. C. (1989). "Aggregation in stochastic ecosystem models." Ecological Modelling **44**: 153-164.

Gaus, G. F. (2006). The Evolution of Society and Mind: Hayek's System of Ideas. The Cambridge Companion to Hayek. E. Feser. Cambridge, U.K., Cambridge University Press**:** 232-258.

Gaylord, R. J. and L. J. D'Andria (1998). Simulating Society: A Mathematica Toolkit for Modeling Socioeconomic Behavior. New York, N.Y., Telos (Springer-Verlag).

Geanakoplos, J., R. L. Axtell, J. D. Farmer, P. Howitt, B. Conlee, J. Goldstein, M. Hendrey, N. M. Palmer and C.-Y. Yang (2012). "Getting at Systemic Risk via an Agent-Based Model of the Housing Market." American Economic Review: Papers and Proceedings **102**(3): 53-58.

Gemann, T. C., K. Kadau, I. M. Longini Jr. and C. A. Macken (2006). "Mitigation strategies for pandemic influenza in the United States." Proc Natl Acad Sci U S A **103**(15): 5935-5940.

Gigerenzer, G. (2000). Adaptive Thinking: Rationality in the Real World. New York, N.Y., Oxford University Press.

Gigerenzer, G. and R. Selten, Eds. (2001). Bounded Rationality: The Adaptive Toolbox. Cambridge, Mass., MIT Press.

Gigerenzer, G., P. M. Todd and A. R. Group (1999). Simple Heuristics That Make Us Smart. New York, Oxford University Press.

Gini, C. (1927). "The Scientific Basis of Fascism." Political Science Quarterly **42**(1): 99-115.
Gollier, C. and R. Zeckhauser (2005). "Aggregation of Heterogeneous Time Preferences." Journal of Political Economy **113**(4): 878-896.
Gorton, G. B. (2010). Slapped by the Invisible Hand: The Panic of 2007. New York, New York, Oxford University Press.
--- (2012). Misunderstanding Financial Crises: Why We Don't See Them Coming. New York, N.Y., Oxford University Press.
Guckenheimer, J. and P. Holmes (1983). Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. New York, N.Y., Springer-Verlag.
Gueron, S. and S. A. Levin (1993). "Self-Organization of Front Patterns in Large Wildebeest Herds." Journal of Theoretical Biology **165**: 541-552.
Gueron, S., S. A. Levin and D. I. Rubenstein (1993). The Dynamics of Mammalian Herds: From Individuals to Aggregations. Working paper, Department of Ecology and Evolutionary Biology. Princeton, N.J., Princeton University Press.
Guerrero, O. A. and R. L. Axtell (2013). "Employment Growth through Labor Flow Networks." PLoS ONE **8**(5): e60808.
Hahn, F. H. (1962). "On the Stability of Pure Exchange Equilibrium." International Economic Review **3**(2): 206-213.
Haken, H. (1987). Synergetics: An Approach to Self Organization. Self-Organizing Systems: The Emergence of Order. F. E. Yates. Berlin, Plenum Press.
Hayek, F. A. v. (1945). "The Use of Knowledge in Society." American Economic Review **35**(4): 519-530.
--- (1955). The Counter-Revolution of Science: Studies in the Abuse of Reason. London, The Free Press.
--- (1967). Studies in Philosophy, Politics and Economics. Chicago, Ill., University of Chicago Press.
--- (1974). The Pretence of Knowledge. Stockholm, Sweden, The Nobel Foundation.
Hebb, D. O. (1949). The Organization of Behavior. New York, N.Y., John Wiley & Sons.
Heilbroner, R. L. (1980). The Worldly Philosophers. New York, N.Y., Simon and Schuster.
Ho, T. H., C. Camerer and J. K. Chong (2004). The economics of learning models: A self-tuning theory of learning in games.
Hoffman, J. and C. Johnson (2007). Computational Turbulent Incompressible Flow. New York, N.Y., Springer.
Holland, J. H. (1976). Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence. Ann Arbor, Mich., University of Michigan Press.
--- (1993). The Effects of Labels (Tags) on Social Interactions. Working paper. Santa Fe, N.M., Santa Fe Institute.
--- (1995). Hidden Order: How Adaptation Builds Complexity. New York, N.Y., Perseus Press.
--- (1998). Emergence: From Chaos to Order. Reading, Mass., Perseus.
--- (2012). Signals and Boundaries: Building Blocks for Complex Adaptive Systems. Cambridge, Mass., MIT Press.
Ijiri, Y. and H. A. Simon (1964). "Business Firm Growth and Size." American Economic Review **54**: 77-89.
--- (1967). "A Model of Business Firm Growth." Econometrica **35**(2): 348-355.
--- (1977). Skew Distributions and the Sizes of Business Firms. New York, N.Y., North-Holland.
Iwasa, Y., V. Andreasen and S. A. Levin (1987). "Aggregation in model ecosystems. I. Perfect aggregation." Ecological Modelling **37**: 287-302.

Iwasa, Y., S. A. Levin and V. Andreasen (1989). "Aggregation in model ecosystems. II. Approximate aggregation." IMA Journal of Mathematics Applied in Medicine and Biology **6**: 1-23.
Jennings, N. R., K. Sycara and M. Wooldridge (1998). "A Roadmap of Agent Research and Development." Autonomous Agents and Multi-Agent Systems **1**(1): 7-38.
Johnson, N. F., P. Jefferies and P. M. Hui (2003). Financial Market Complexity: What Physics Can Tell Us About Market Behavior. New York, N.Y., Oxford University Press.
Kahneman, D., J. L. Knetsch and R. H. Thaler (1986). Fairness and the Assumptions of Economics. Rational Choice: The Contrast between Economics and Psychology. R. M. Hogarth and M. W. Reder. Chicago, Ill., University of Chicago Press.
Kahneman, D. and A. Tversky (1979). "Prospect Theory: An Analysis of Decisions Under Risk." Econometrica **47**: 313-327.
Kaldor, N. (1972). "The Irrelevance of Equilibrium Economics." Economic Journal **82**(328): 1237-1255.
--- (1985). Economics without Equilibrium. Cardiff, U.K., University College Cardiff Press.
Karr, J. R., J. C. Sanghvi, D. N. Macklin, M. V. Gutschow, J. M. Jacobs, B. J. Bolival, N. Assad-Garcia, J. I. Glass and M. W. Covert (2012). "A Whole-Cell Computational Model Predicts Phenotype from Genotype." Cell **150**: 389-401.
Kauffman, S. (1995). At Home in the Universe: The Search for Laws of Self-Organization and Complexity. New York, N.Y., Oxford University Press.
Keynes, J. M. (1936). The General Theory of Employment, Interest, and Money.
Kirman, A. P. (1989). "The Intrinsic Limits of Modern Economic Theory: The Emperor Has No Clothes." Economic Journal **99**(395): 126-139.
Klamer, A. (1984). The New Classical Macroeconomics: Conversations with New Classical Economists and their Opponents. Brighton, U.K., Wheatsheaf.
Koppl, R. (2009). Complexity and Austrian economics. Handbook of Complexity Research. J. Rosser, J. Barkley. Cheltenham, U.K., Edward Elgar.
Laughlin, R. B. and D. Pines (2000). "The Theory of Everything." Proc Natl Acad Sci U S A **97**(1): 28-31.
Lewars, E. G. (2011). Computational Chemistry: Introduction to the Theory and Applications of Molecular and Quantum Mechanics. New York, N.Y., Springer.
Lomuscio, A. and M. Sergot (2002). On Multi-agent Systems Specification vis Deontic Logic. Intelligent Agents VIII. J.-J. C. Meyer and M. Tambe. Berlin, Springer-Verlag. **2333:** 86-99.
Lorenz, E. N. (1963). "Deterministic Nonperiodic Flow." Journal of Atmospheric Sciences **20**: 130-141.
Luce, R. D. and H. Raiffa (1957). Games and Decisions: Introduction and Critical Survey. New York, N.Y., John Wiley & Sons.
Luckyanov, N. K., Y. M. Svirezhev and O. V. Voronkova (1983). "Aggregation of variables in simulation models of water resources." Ecological Modeling **18**: 235-240.
Lux, T. (1998). "The Socioeconomic Dynamics of Speculative Markets: Interacting Agents, Chaos and the Fat Tails of Return Distributions." Journal of Economic Behavior and Organization **33**: 143-165.
Lux, T. and M. Marchesi (1999). "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market." Nature **397**: 498-500.
Mackay, C. (1841). Extraordinary Popular Delusions and the Madness of Crowds. London, U.K., Richard Bentley.
Mackenzie, L. W. (1954). "On Equilibrium in Graham's Model of World Trade and Other Competitive Systems." Econometrica **22**(2): 147-161.
Mandelbrot, B. (1983). The Fractal Geometry of Nature. New York, N.Y., W.H. Freeman.

--- (1997). Fractals and Scaling in Finance: Discontinuity, Concentration and Risk. New York, N.Y., Springer-Verlag.
Mantegna, R. N. and H. E. Stanley (2000). An Introduction to Econophysics: Correlations and Complexity in Finance. New York, N.Y., Cambridge University Press.
Markram, H. (2006). "The Blue Brain Project." Nature Reviews Neuroscience **7**: 153-160.
--- (2012). "A Countdown to a Digital Simulation of Every Last Neuron in the Human Brain." Scientific American.
May, R. M. (1974). Stability and Complexity in Model Ecosystems. Princeton, N.J., Princeton University Press.
McClelland, J. L. and D. E. Rumelhart (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Cambridge, Mass., MIT Press.
Menger, C. (1985 [1883]). Investigations into the Method of the Social Sciences. Grover City, Penn., Libertarian Press, Inc.
Michalewicz, Z. and D. B. Fogel (2000). How to Solve It: Modern Heuristics. New York, N.Y., Springer-Verlag.
Miller, J. H. and S. E. Page (2007). Complex Adaptive Systems: An Introduction to Computational Models of Social Life. Princeton, N.J., Princeton University Press.
Minsky, M. and S. Papert (1969). Perceptrons. Cambridge, Mass., MIT Press.
Montroll, E. W. and W. W. Badger (1974). Introduction to Quantitative Aspects of Social Phenomena. New York, N.Y., Gordon and Breach Science Publishers.
Morowitz, H. J. (1998). "Emergence and Equilibrium." Complexity **4**(6): 12-13.
--- (2002). The Emergence of Everything: How the World Became Complex. New York, N.Y., Oxford University Press.
Nagel, K. and C. Barrett (1997). "Using Microsimulation Feedback for Trip Adaptation for Realistic Traffic in Dallas." International Journal of Modern Physics C **8**(3): 505-525.
Nagel, K., R. Beckman and C. L. Barrett (1998). TRANSIMS for Transportation Planning. Technical Report. Los Alamos, N.M., Los Alamos National Laboratory.
Nagel, K. and M. Paczuski (1995). "Emergent Traffic Jams." Physical Review E **51**: 2909.
Nagel, K. and S. Rasmussen (1994). Traffic at the Edge of Chaos. Artificial Life IV. R. A. Brooks and P. Maes. Cambridge, Mass., MIT Press: 224-235.
Negishi, T. (1961). "On the Formation of Prices." International Economic Review **2**(1): 122-126.
Newell, A. and H. A. Simon (1972). Human Problem Solving. Englewood Cliffs, N.J., Prentice-Hall.
Nilsson, M. N. and O. Görnerup (2008). A Dual Eigenvector Condition for Strong Lumpability of Markov Chains. Santa Fe Institute working paper series. Santa Fe, New Mexico.
Nozick, R. (1974). Anarchy, State and Utopia. New York, N.Y., Basic Books.
--- (1994). "Invisible Hand Explanations." American Economic Review **84**(2): 314-318.
Perelson, A. S., Ed. (1988). Theoretical Immunology. Redwood City, Calif., Addison-Wesley.
Perline, R., R. Axtell and D. Teitelbaum (2006). Volatility and Asymmetry of Small Firm Growth Rates Over Increasing Time Frames. SBA Research Reports. Washington, D.C.
Polanyi, M. (1974). Personal Knowledge: Towards a Post-Critical Philosophy. Chicago, Ill., University of Chicago Press.
Railsback, S. F. and V. Grimm (2012). Agent-Based and Individual-Based Modeling: A Practical Introduction. Princeton, N.J., Princeton University Press.
Resnick, M. (1994). Turtles, Termites and Traffic Jams: Explorations in Massively Parallel Microworlds. Cambridge, Mass., MIT Press.
Rosser, J., J. Barkley (2009). How Complex are the Austrians. Harrisonburg, Virginia: 1-22.

Roth, A. E. and I. Erev (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." Games and Economic Behavior **8**: 164-212.
Rothschild, E. (1994). "Adam Smith and the Invisible Hand." American Economic Review **84**(2): 319-322.
Rumelhart, D. E. and J. L. McClelland (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition. Cambridge, Mass., MIT Press.
Russell, S. J. and P. Norvig (2002). Artificial Intelligence: A Modern Approach. Englewood Cliffs, N.J., Prentice Hall.
Scarf, H. E. (1960). "Some Examples of Global Instability of the Competitive Equilibrium." International Economic Review **1**: 157-172.
Schelling, T. C. (1971). "Dynamic Models of Segregation." Journal of Mathematical Sociology **1**: 143-186.
--- (1971). "On the Ecology of Micro-Motives." Public Interest **25**: 59-98.
Seeley, T. D. (1995). The Wisdom of the Hive: The Social Psychology of Honeybee Colonies. Princeton, N.J., Princeton University Press.
Seuss, D. (1961). The Sneetches and Other Stories. New York, N.Y., Random House.
Shoham, Y., R. Powers and T. Grenager (2004). Multi-agent reinforcement learning: A critical survey. AAAI Fall Symposium on Artifical Multi-Agent Learning.
Simon, H. A. (1955). "A Behavioral Model of Rational Choice." Quarterly Journal of Economics **69**(1): 99-118.
--- (1955). "On a Class of Skew Distribution Functions." Biometrika **42**: 425-440.
--- (1956). "Rational Choice and the Structure of the Environment." Psychological Review **63**(2): 129-138.
--- (1957). Models of Man: Social and Rational. New York, N.Y., John Wiley & Sons, Inc.
--- (1996). The Sciences of the Artificial. Cambridge, Mass., MIT Press.
Simon, H. A. and C. Bonini (1958). "The Size Distribution of Business Firms." American Economic Review **48**(4): 607-617.
Slovic, P., B. Fischoff and S. Lichtenstein (1982). Facts versus Fears: Understanding Perceived Risk. Judgement Under Uncertainty: Heuristics and Biases. D. Kahneman, P. Slovic and A. Tversky. New York, N.Y., Cambridge University Press.
Stokey, N. L. and R. E. Lucas, Jr. (1989). Recursive Economic Dynamics. Cambridge, Mass., Harvard University Press.
Strogatz, S. H. (2003). Sync: The Emerging Science of Spontaneous Order. New York, N.Y., Theia.
Sugden, R. (1989). "Spontaneous Order." Journal of Economic Perspectives **3**(4): 85-97.
Tversky, A. and D. Kahneman (1981). "The Framing of Decisions and the Rationality of Choice." Science **211**: 453-458.
--- (1986). Rational Choice and the Framing of Decisions. Rational Choice: The Contrast between Economics and Psychology. R. M. Hogarth and M. W. Reder. Chicago, Ill., University of Chicago Press.
--- (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty." Journal of Risk and Uncertainty **5**(297-323).
Uzawa, H. (1962). "On the Stability of Edgeworth's Barter Process." International Economic Review **3**(2): 218-232.
Vaughn, K. (1999). "Hayek's Theory of the Market Order as an Instance of the Theory of Complex Adaptive Systems." Journal de Economistes et des Etudes Humaines **9**: 241-256.
Veblen, T. (1918). The Higher Learning in America: A Memorandum On the Conduct of Universities by Business Men. New York, N.Y., B.W. Huebsch.

Vriend, N. J. (2002). "Was Hayek an Ace?" Southern Economic Journal **68**(4): 811-840.
Waldrop, M. M. (1992). Complexity: The Emerging Science at the Edge of Order and Chaos. New York, N.Y., Simon & Schuster.
Weintraub, E. R. (2011). "Lionel W. McKenzie and the Proof of the Existence of a Competitive Equilibrium." Journal of Economic Perspectives **25**(2): 199-215.
Wooldridge, M. (2002). An Introduction to Multi-Agent Systems. West Sussex, England, John Wiley & Sons.