Complexity economics:
a different framework for economic thought

W. Brian Arthur [1]

March 12, 2013

Over the last twenty-five years, a different approach to economics has been slowly birthing, and slowly growing—*complexity economics*. Complexity economics holds that the economy is not necessarily in equilibrium, that computation as well as mathematics is useful in economics, that increasing as well as diminishing returns may be present in an economic situation, and that the economy is not something given and existing but forms from a constantly developing set of institutions, arrangements, and technological innovations. The approach got its start largely at the Santa Fe Institute in the late 1980s but now has many practitioners,[2] and it raises several questions. What does this different way of thinking about the economy offer? How exactly does it work and where does it fit in? Will it replace neoclassical economics, or be subsumed into neoclassical economics? And under what logic, if any, does it operate?

My purpose in this paper is to answer these questions, especially the last one. In doing so I will not attempt to provide a survey or guided tour, rather I want to provide a framework—a coherent logic—for thinking about this new approach. I will argue from first principles and will build from two earlier essays of mine (Arthur 1999, 2006) as well as the work of many other people to illustrate the key points.[3]

I will argue that this new approach is not just an extension of standard economics, nor does it consist of adding agent-based behavior to standard models. It is a different way of seeing the economy. It gives a different view, one where actions and strategies constantly evolve, where time becomes important, where structures constantly form and re-form, where phenomena appear that are not visible to standard equilibrium analysis, and where a meso-layer between the micro and the macro becomes important. This view, in other words, gives us a world closer to that of political economy than to neoclassical theory, a world that is organic, evolutionary, and historically-contingent.

---

[2] See *The Economy as an Evolving Complex System* volumes edited by: Arrow, Anderson and Pines (1988); Arthur, Durlauf and Lane (1997); and Blume and Durlauf (2005). For history of the ideas see Fontana (2010), Arthur (2010b), and the popular accounts of Waldrop (1992) and Beinhocker (2006). Variants of complexity economics include generative economics, interactive-agent economics, agent-based computational economics, (see Epstein 2006a; Miller and Page, 2007; Tesfatsion and Judd, 2006).

[3] For other essays on this general approach see Axtell (2007), Colander (2000, 2012), Epstein (2006), Farmer (2012), Judd (2006), Kirman (2011), Rosser (1999), and Tesfatsion (2006). The term "complexity economics" was first used in Arthur (1999).

The Economy and Complexity

Let me begin with the economy itself. The economy is a vast and complicated set of arrangements and actions wherein agents—consumers, firms, banks, investors, government agencies—buy and sell, speculate, trade, oversee, bring products into being, offer services, invest in companies, strategize, explore, forecast, compete, learn, innovate, and adapt. In modern parlance we would say it is a massively parallel system of concurrent behavior. And from all this concurrent behavior markets form, prices form, trading arrangements form, institutions and industries form. Aggregate patterns form.

One of the earliest insights of economics—it certainly goes back to Smith—is that these aggregate patterns form from individual behavior, and individual behavior in turn responds to these aggregate patterns: there is a recursive loop. It is this recursive loop that connects with complexity. Complexity is not a theory but a movement in the sciences that studies how the interacting elements in a system create overall patterns, and how these overall patterns in turn cause the interacting elements to change or adapt. It might study how individual cars together act to form patterns in traffic, and how these patterns in turn cause the cars to alter their position. Complexity is about formation—the formation of structures—and how this formation affects the objects causing it.

To look at the economy, or areas within the economy, from a complexity viewpoint then would mean asking how it evolves, and this means examining in detail how individual agents' behaviors together form some outcome and how this might in turn alter their behavior as a result. Complexity in other words asks how individual behaviors might *react to* the pattern they together create, and how that pattern would alter itself as a result. This is often a difficult question; we are asking how a process is created from the purposed actions of multiple agents. And so economics early in its history took a simpler approach, one more amenable to mathematical analysis. It asked not how agents' behaviors would *react to* the aggregate patterns these created, but what behaviors (actions, strategies, expectations) would be upheld by—would be *consistent with*—the aggregate patterns these caused. It asked in other words what patterns would call for no changes in micro-behavior, and would therefore be in stasis, or equilibrium. (General equilibrium theory thus asked what prices and quantities of goods produced and consumed would be consistent with—would pose no incentives for change to—the overall pattern of prices and quantities in the economy's markets. Classical game theory asked what strategies, moves, or allocations would be consistent with—would be the best course of action for an agent (under some criterion)—given the strategies, moves, allocations his rivals might choose. And rational expectations economics asked what expectations would be consistent with—would on average be validated by—the outcomes these expectations together created.)

This equilibrium shortcut was a natural way to examine patterns in the economy and render them open to mathematical analysis. It was an understandable—even proper—way to push economics forward. And it achieved a great deal. Its central construct, general equilibrium theory, is not just mathematically elegant; in modeling the economy it re-composes it in our minds, gives us a way to picture it, a way to comprehend the economy in its wholeness. This is extremely valuable, and the same can be said for other equilibrium modelings: of the theory of the firm, of international trade, of financial markets.

But there has been a price for this equilibrium finesse. Economists have objected to it—to the neoclassical construction it has brought about—on the grounds that it posits an idealized, rationalized world that distorts reality, one whose underlying assumptions are

often chosen for analytical convenience.[4] I share these objections. Like many economists I admire the beauty of the neoclassical economy; but for me the construct is too pure, too brittle—too bled of reality. It lives in a Platonic world of order, stasis, knowableness, and perfection. Absent from it is the ambiguous, the messy, the real.

Good economists of course have always harbored a richer view of the economy than this (Colander and Kupers, 2012; Louça, 2010), so perhaps we could stick with equilibrium as the basis of our thinking, allowing that experience and intuition can fill out the realities. But this still is not satisfactory. If we assume equilibrium we place a very strong filter on what we can see in the economy. Under equilibrium by definition there is no scope for improvement or further adjustment, no scope for exploration, no scope for creation, no scope for transitory phenomena, so anything in the economy that takes adjustment—adaptation, innovation, structural change, history itself—must be bypassed or dropped from theory. The result may be a beautiful structure, but it is one that lacks authenticity, aliveness, and creation.

What if economics allowed the wider possibility and asked how agents in the economy might react to the patterns they together create? Would this make a difference? What would we see then?

Endogenously Generated Nonequilibrium

The first thing to observe is that in asking "how agents might react to," we are implicitly assuming nonequilibrium, for if novel reactions are possible they will alter the outcome, so by definition it cannot be an equilibrium. A well-trained economist might object to this assumption of nonequilibrium; standard doctrine holds that nonequilibrium cannot be important in the economy. "[P]ositions of unstable equilibrium," said Samuelson (1983), "even if they exist, are transient, non-persistent states. … How many times has the reader seen an egg standing on its end?"[5]

Equilibrium, we are assured, is the natural state of the economy.

I want to argue that this is not the case, emphatically not the case, that nonequilibrium is the natural state of the economy, and therefore the economy is always open to reaction. This isn't merely because of outside shocks or external influences, but because nonequilibrium arises endogenously in the economy. There are two main reasons for this. One is fundamental (or Knightian) uncertainty, the other is technological innovation. Let me take each in turn.

First, fundamental uncertainty. All problems of choice in the economy involve something that takes place in the future, perhaps almost immediately, perhaps at some distance of time. Therefore they involve some degree of not knowing. In some cases agents are well informed, or can put realistic probability distributions over events that might happen; but in many other cases—in fact in most cases—they have no basis to do this, they

---

[4]Blaug (2003), Bronk (2009, 2011), Cassidy (2009), Colander et al. (2009), Davis (2007), Farmer and Geanakoplos, (2008), Kirman (2010), Koppl and Luther (2010), Krugman (2009), Mirowski (2002), Simpson (2002).

[5] Walras expressed a similar thought in a 1909 conversation with Schumpeter, "life is essentially passive and merely adapts itself to the natural and social influences which may be acting on it, so that the theory of a stationary process constitutes really the whole of theoretical economics …." (Tabb, 1999; Reisman, 2004).

simply do not know.[6] I may be choosing to put venture capital into a new technology, but my startup may simply not know how well the technology will work, how the public will receive it, how the government will choose to regulate it, or who will enter the space with a competing product. I must make a move but I have genuine not-knowingness—fundamental uncertainty. There is no "optimal" move. Things worsen when other agents are involved; such uncertainly then becomes self-reinforcing. If I cannot know exactly what the situation is, I can take it that other agents cannot know either. Not only will I have to form subjective beliefs, but I will have to form subjective beliefs about subjective beliefs. And other agents must do the same. Uncertainty engenders further uncertainty. [7]

This observation of course is not new. Other economists, Shackle in particular (1955, 1992), have written much about this. But it has an important consequence for theorizing. To the degree that outcomes are unknowable, the decision problems they pose are not well-defined. It follows that rationality—pure deductive rationality—is not well-defined either, for the simple reason that there cannot be a logical solution to a problem that is not logically defined. It follows that in such situations deductive rationality is not just a bad assumption; it cannot exist. There might be intelligent behavior, there might be sensible behavior, there might be farsighted behavior, but rigorously speaking there can not be deductively rational behavior. Therefore we cannot assume it.

None of this means that people cannot proceed in the economy, or that they do not choose to act. Behavioral economics tells us that often the context determines how people decide, and certainly we can use its findings. And cognitive science tells us that if a decision is important, people may stand back from the situation and attempt to make sense out of it by surmising, making guesses, using past knowledge and experience. They use their imaginations to try to come up with some picture of the future and proceed on this (Bronk, 2009, 2013). Indeed, as Shackle (1992) puts it, "The future is imagined by each man for himself and this process of the imagination is a vital part of the process of decision." One way to model this is to suppose economic agents form individual beliefs (possibly several) or hypotheses—internal models—about the situation they are in and continually update these, which means they constantly adapt or discard and replace the actions or strategies based on these as they explore.[8] They proceed in other words by induction (Holland et al., 1986; Sargent, 1993; Arthur, 1994).[9]

This ongoing materialization of exploratory actions causes an always-present Brownian motion within the economy. The economy is permanently in disruptive motion as agents explore, learn, and adapt. These disruptions, as we will see, can get magnified into larger phenomena.

---

[6] As Keynes (1937) puts it: "the prospect of a European war … the price of copper … the rate of interest twenty years hence…. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know."

[7] Soros (1987) calls this the principle of *reflexivity*.

[8] A standard objection is that allowing agents to reason non-deductively admits arbitrariness. What prevents such beliefs or behaviors from being chosen *ad-hoc* to yield some favored outcome? Certainly this is possible, but that doesn't justify retreating to unrealistic "rational" models of behavior. The idea is not to assume behavior that makes analysis simple, but behavior that makes models realistic.

[9] Calling this "bounded rationality" is a misnomer. It implies that agents do not use all reasoning powers at their disposal, which under uncertainly may often be false.

The other driver of disruption is technological change. About a hundred years ago, Schumpeter (1912) famously pointed out that there is "a source of energy within the economic system which would of itself disrupt any equilibrium that might be attained." That source was "new combinations of productive means." (Nowadays we would say new combinations of technology.) Economics does not deny this, but incorporates it by allowing that from time to time its equilibria must adjust to such outside changes.

But this technology force is more disruptive than Schumpeter allowed. Novel technologies call forth further novel technologies: when computers arrive, they call forth or "demand" the further technologies of data storage, computer languages, computational algorithms, and solid-state switching devices. And novel technologies make possible other novel technologies: when the vacuum tube arrives, it makes possible or "supplies" the further technologies of radio transmission and receiving, broadcasting, relay circuits, early computation, and radar. And these novel technologies in turn demand and supply yet further technologies. It follows that a novel technology is not just a one-time disruption to equilibrium, it is a permanent ongoing generator and demander of further technologies that themselves generate and demand still further technologies (Arthur, 2009). Notice again the self-reinforcing nature of this process. The result is not occasional disruption but ongoing waves of disruption causing disruptions, acting in parallel across the economy and at all scales within the economy. Technology change breeds further change endogenously and continually, and this throws the economy into a permanent state of disruption.

Technological disruption acts on a somewhat slower timescale than the Brownian motion of uncertainty. But if anything it causes larger upheavals. And by itself it induces further uncertainty: businesses and industries simply do not know what technologies will enter their space next. Both uncertainty and technology then give us an economy where agents have no determinate means to make decisions.

A picture is now emerging of the economy different from the standard equilibrium one. To the degree that uncertainty and technological changes are present in the economy—and certainly both are pervasive at all levels—agents must explore their way forward, must "learn" about the decision problem they are in, must respond to the opportunities confronting them. We are in a world where beliefs, strategies, and actions of agents are being "tested" for survival within a situation or outcome or "ecology" that these beliefs, strategies and actions together create. Further, and more subtly, these very explorations alter the economy itself and the situation agents encounter. So agents are not just reacting to a problem they are trying to make sense of; their very actions in doing so collectively re-form the current outcome, which requires them to adjust afresh. We are, in other words, in a world of complexity, a complexity closely associated with nonequilibrium.

Theorizing under Nonequilibrium

Where does this leave us? If the economy is large and constantly aboil with activity, then we would seem to be dealing here (to borrow a phrase from Schumpeter, 1954) with "a chaos that is not in analytical control." Faced with this prospect in the past, economics has metaphorically thrown up its hands and backed away. But what if we don't do this, what if we stand our ground and take nonequilibrium seriously, how then can we proceed? Can we say anything useful? What would we see? And above all, what would it mean to do theory under nonequilibrium?

Certainly, many parts of the economy could be still be treated as approximately at equilibrium, and standard theory would still be valid here. And other parts could be treated

as temporarily diverging from strong attracting states, and we could study convergence here. But this would still be seeing the economy as a well-balanced machine temporarily prone to getting out of adjustment; and that neither gets us to the heart of seeing how the economy behaves out of equilibrium nor captures the creative side of nonequilibrium.

A better way forward is to observe that in the economy, current circumstances form the conditions that will determine what comes next. The economy is a system whose elements are constantly updating their behavior based on the present situation.[10] To state this in another way, formally, we can say that the economy is an ongoing *computation*—a vast, distributed, massively parallel, stochastic one.[11] Viewed this way, the economy becomes a system that evolves procedurally in a series of events; it becomes algorithmic.

There is a danger that seeing the economy this way is merely bowing to a current fashion in science, but the idea allows me to make an important point. Suppose for a moment that we—or better, Laplace or "God"—know the algorithm[12] behind the computation (the large but finite set of detailed mechanisms by which the economy, or the part of it that interests us, makes its next move). A fundamental theorem in computation (Turing, 1936) tells us that in general (if we choose an algorithm randomly) there is no way—no systematic analytical method—to tell in advance whether that algorithm or computer program will halt (as opposed to going on forever, or cycling). Since we could arrange that an algorithm halt if its output fulfilled some particular set of mathematical conditions or reached a given "solution," in general we cannot tell that that will be the case either. In other words there is no analytical method to decide in advance what a given algorithm will do.[13] All we can do is follow the computation and see what it brings. Of course, with simple algorithms we can often see they will settle down to a given outcome. But algorithms don't have to be particularly complicated before we cannot decide their outcomes (Wolfram, 2002).

So we need to be cautious. For highly interconnected systems, equilibrium and closed-form solutions are not the default outcomes; if they exist they require justification. And computation for such systems should not be regarded as the avoidance of analytical thinking; rigorously speaking, it may be completely necessary. We can often do much useful pre-analysis of the qualitative properties of nonequilibrium systems, and understand the mechanisms behind these; still, in general the only precise way to study their outcomes is by computation itself.

Of course the algorithm behind the actual economy is not randomly chosen, it is highly structured, so it may be that the actual economy's "computations" always have simple outcomes. Or it may equally be that the economy's computations are always unordered and amorphous. Usually in the parts of the economy we study, neither is the case. Often, especially when there are strong countervailing forces at work, we see large structures— regions of attraction that correspond loosely to equilibria. And within these (or in their absence) we also see mechanisms that cause phenomena or sub-patterns or substructures to appear and disappear randomly from time to time. To give a physical analogy, consider the sun. From afar it appears to be a large gaseous ball in uniform spherical equilibrium. But

---

[10] Current circumstances would of course include relevant past history or memory of past history.

[11] Modern computational thinking sees computation as ongoing, concurrent (parallel), distributed, and often probabilistic. See the 2010 ACM Ubiquity Symposium *What is Computation?* See also Beinhocker (2011).

[12] Earlier I argued that the economy's future is indeterminate, so strictly speaking the economy is not perfectly algorithmic. Hence for this thought experiment I posit a "God" who can determine how each agent would react in all circumstances.

[13] Including whether it converges (or stays within a given neighborhood of some limit forever).

within this "equilibrium," powerful mechanisms cause dynamic phenomena such as gigantic magnetic loops and arches, coronal holes, X-ray bright spots, and mass plasma ejections moving at up to 2,000 kilometers per second. The gaseous ball indeed displays a loose spherical shape, but it is never at equilibrium. Rather it is seething with activity that disrupts the possibility of equilibrium and builds from earlier disruptions. These phenomena are localized and can act at many scales. And they are transitory or temporal—they appear, disappear, and interact, seemingly randomly in time.

We will find a similar situation frequently in the economy. Theorizing in nonequilibrium then would mean uncovering large attractors at work (if indeed there are any), but also studying other sub-structures or phenomena that might be present for their properties and behavior. We can use carefully-designed computer experiments to do this, often using statistics on the results to isolate phenomena and the mechanisms that cause these. And in many cases we can construct simpler toy models of a phenomenon that capture its essential features and allow us to use mathematics or stochastic theory to study it. The objective, we should remember, is not necessarily to formulate equations or to arrive at necessary conditions. The objective, as it is with all theory, is to obtain general insights.

Let us put some of these ideas together by looking at an actual nonequilibrium study performed computationally. Here is a classic example.

In 1991 Kristian Lindgren constructed a computerized tournament where strategies competed in randomly chosen pairs to play a repeated prisoner's dilemma game. (The details of the prisoner's dilemma needn't concern us; think of this as simply a game played by a specified current set of strategies.) The strategies consisted of instructions for how to move given the opponent's immediate past moves. If strategies did well they replicated and mutated, if they did badly they were removed. Lindgren allowed that strategies could "deepen" by using deeper memory of their opponent's immediate past moves and their own. So in our language we can think of such strategies as "exploring" strategy space: they change and adapt if they are not successful. Lindgren found that at the start of his tournament, simple strategies such as Tit-for-Tat dominated, but over time, deeper strategies appeared that exploited the simple ones. In time, further deeper strategies emerged to take advantage of these with periods of relative stasis alternating with dynamic instability (Fig. 1).

The dynamics are simple enough that Lindgren could write them as stochastic equations, yet these give far from a full picture; we really need computation to see what is going on. What emerges computationally is an *ecology*—an ecology of strategies, each attempting to exploit and survive within an environment created by itself and other strategies attempting to exploit and survive. This ecology is a miniature biosphere where novel species (strategies) continually appear, exploit the environment created by existing species, and do not survive if they fail. Notice that evolution has entered, but it hasn't been brought in from outside, it has arisen in the natural tendency of strategies to compete for survival. The point is general in this type of economics. What constitutes a "solution" is typically an ecology where strategies, or actions, or beliefs compete; an ecology that may not settle down, that has its own characteristic properties and can be studied qualitatively and statistically. [14]

In Lindgren's study, the outcome differs from one run of the computation to another. In many runs an evolutionary stable strategy appears, a complicated one that relies on four

---

[14] In the well-known El Farol problem (Arthur, 1994) an ecology of ever-changing individual forecasts emerges, along with overall a equilibrium attractor state. Metaphorically the individual trees change, but the shape of the forest persists.

periods of memory of past actions. In other runs the system continues to evolve. In some runs we see the quick emergence of complicated strategies, in others these appear later on. And yet there are constants: phenomena such as coexistence among strategies, exploitation, the spontaneous emergence of mutualism, sudden collapses, periods of stasis and unstable change. The picture resembles paleozoology more than anything else.
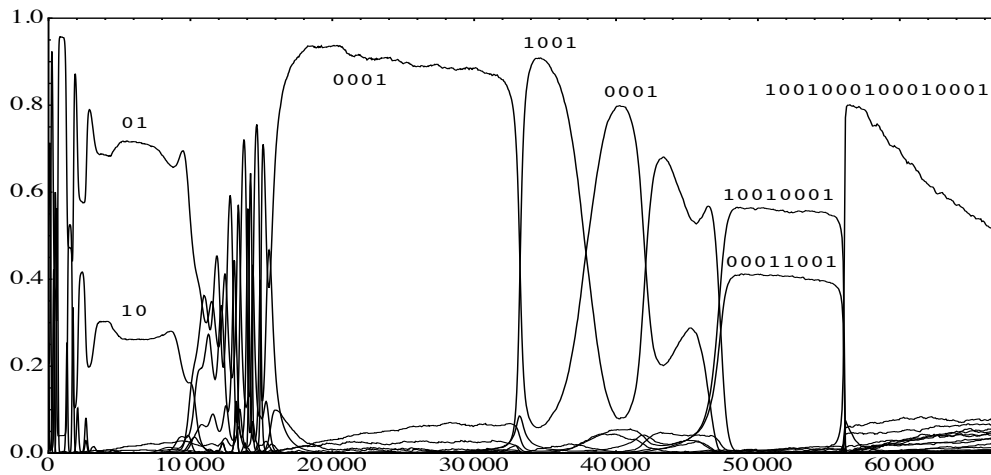

Figure 1. Strategies in Lindgren's computerized tournament. The horizontal axis denotes time, the vertical axis numbers using a particular strategy, the labels code for the memory-depth of strategies.

I have put forward Lindgren's study as an example of doing nonequilibrium economics and the reader may be wondering how the study of such computer-based worlds can qualify as economics, or what relationship this might have to doing theory. My answer is that theory does not consist of mathematics. Mathematics is a technique, a tool, albeit a sophisticated one. Theory is something different. Theory lies in the discovery, understanding, and explaining of phenomena present in the world. Mathematics facilitates this—enormously— but then so does computation. Naturally, there is a difference. Working with equations allows us to follow an argument step by step and reveals conditions a solution must adhere to, whereas computation does not.[15] But computation—and this more than compensates—allows us to see phenomena that equilibrium mathematics does not. It allows us to rerun results under different conditions, exploring when structures appear and don't appear, isolating underlying mechanisms, and simplifying again and again to extract the bones of a phenomenon. Computation in other words is an aid to thought, and it joins earlier aids in economics—algebra, calculus, statistics, topology, stochastic processes—each of which was resisted in its time. The computer is an exploratory lab for economics, and used skillfully, a powerful generator for theory.[16]

All this suggests a way forward for our nonequilibrium way of looking at the economy. We can see the economy, or the parts of it that interest us, as the ever-changing outcome of agents' strategies, forecasts, and behaviors. And we can investigate these parts, and also

---

[15] Note that we can always rewrite any algorithmic model in equation form (any computation by a Turing machine can be represented in equation form) so that, rigorously speaking, computation-based analysis is as mathematical as standard analysis. See Epstein (2006).

[16] For computation's role in theorizing in mathematics, physics, biology and earth-sciences, see Robertson (2003). See also Bailey (2010) and Chaitin (2006).

classic problems within economics—intergenerational transfers, asset pricing, international trade, financial transactions, banking—by constructing models where responses are specified not just at equilibrium but in all circumstances. Sometimes our models will be amenable to mathematical analysis, sometimes only to computation, sometimes to both. What we can seek is not just equilibrium conditions, but understandings of the formation of outcomes and their further unfolding, and of any dynamic phenomena that appear.

Phenomena and How they Arise

What dynamic phenomena then appear under nonequilibrium? And how do these, and nonequilibrium, connect with complexity? I will take these two questions in succession. To look at what patterns or structures might appear in the economy under nonequilibrium, we can begin by looking at the difference the filter of equilibrium makes to the patterns we see. To set ideas, consider a simple model of something slightly outside the economy, traffic flow.

A typical model would acknowledge that at close separation from cars in front, cars lower their speed, and at wide separation they raise it. A given high density of traffic of N cars per mile would imply a certain average separation, and cars would slow or accelerate to a speed that corresponds. Trivially, an equilibrium speed emerges, and if we were restricting solutions to equilibrium that is all we would see. But in practice at high density, a nonequilibrium phenomenon occurs. Some car may slow down—its driver may lose concentration or get distracted—and this might cause cars behind to slow down. This immediately compresses the flow, which causes further slowing of the cars behind. The compression propagates backwards, traffic backs up, and a jam emerges. In due course the jam clears. But notice three things. The phenomenon's onset is spontaneous; each instance of it is unique in time of appearance, length of propagation, and time of clearing. It is therefore not easily captured by closed-form solutions, but best studied by probabilistic or statistical methods. Second, the phenomenon is *temporal*, it emerges or happens within time, and cannot appear if we insist on equilibrium.[17] And third, the phenomenon occurs neither at the micro-level (individual car level) nor at the macro-level (overall flow on the road) but at a level in between—the *meso-level*.

What about the economy more generally? If we are willing to take away the equilibrium filter, what phenomena might we see there and how will these operate? I will mention three.

The first is *self-reinforcing asset-price changes,* or in the vernacular, bubbles and crashes. To see how these are generated consider the Santa Fe artificial stock market (Palmer et al., 1994; Arthur et al., 1997). In this computer-based model the "investors" are artificially-intelligent computer programs, who for the reasons given earlier, cannot simply assume or deduce a given "rational" forecasting model, but must individually discover expectations (forecasting models) that work well. The investors randomly generate (or discover) their own forecasting methods, try out promising ones, drop those that don't work, and periodically generate new ones to replace them. The stock price forms from their bids and

---

[17] We could of course model this as a stationary stochastic process that includes jams, and artificially call this an "equilibrium" process. Some neoclassical models do this (e.g., Angeletos and La'O, 2011), which would seem to negate my claim that standard economics doesn't handle nonequilibrium. But closer scrutiny shows that such nonequilibrium behavior is always contained within an overall equilibrium wrapper, typically within some overall stochastic process that remains stationary (and hence "in equilibrium"). Such models stretch the neoclassical paradigm by appearing to be "in equilibrium," but at their core are nonequilibrium processes, so I include them as such under the argument here.

offers, and thus ultimately from agents' forecasts. Our market becomes an ecology of forecasting methods that either succeed or are winnowed out, an ecology that perpetually changes as this happens.[18] And we see several phenomena, chief among them, spontaneous bubbles and crashes.

To see how these appear, we can extract a simple version of the mechanism from our experiment. Suppose some of our investors "discover" a class of trading forecast that essentially says "If the price has risen in the last k periods, expect it to increase by x% next period." Suppose also, some investors (they could even be the same investors) "discover" forecasts of the type: "If the current price is more than y times fundamental earnings (or dividend) value, expect it to fall by z%." The first forecasts cause bubble behavior: if the price rises for a while, investors will buy in thus validating it, which may cause a further rise. Eventually this drives the price high enough to trigger the second type of forecast. Investors holding these sell, the price drops, which switches off the upward forecasts, causing other investors to sell too, and a crash ensues. The scale and duration of such disruptions vary, they happen randomly in time, so they cannot be predicted. What *can* be predicted is that such phenomena will occur, and will have certain probability distributions of size and scale.

A second temporal phenomenon is *clustered volatility*. This is the appearance of random periods of low activity followed by periods of high activity. In our artificial market these show up as periods of low and high price volatility. Low volatility reigns when agents' forecasts are working reasonably well mutually; then there is little incentive to change them or the results they produce. High volatility happens when some agent or group of agents "discover" better predictors. This perturbs the overall pattern, so that other investors have to change *their* predictors to readapt, causing further perturbation and further re-adaptation. (This pattern is clearly visible in Lindgren's study, Fig. 1.) The result is a period of intense readjustment or volatility. Such random periods of low volatility alternating with high volatility show up in actual financial market data, where they are called GARCH behavior.

A third phenomenon, more to do with space than with time, we can call *sudden percolation*. When a transmissible change happens somewhere in a network, if the network is sparsely connected the change will sooner or later peter out for lack of onward connections. If the network is densely connected, the change will propagate and continue to propagate. In a network of banks, an individual bank might discover it holds distressed assets. It then comes under pressure to increase its liquidity and calls on its counterparty banks. These in turn come under pressure to increase their liquidity and call on *their* counterparties, and so the distress cascades across the network (Haldane, 2009). Such events can cause serious damage. They peter out in a low-connection network, but propagate—or percolate—for long periods as the degree of connection passes some point and gets large (Watts, 2002).[19]

This last example brings us to a general property. Generally in complex systems, phenomena do not appear until some underlying parameter of the model that depicts the intensity of adjustment or the degree of connection passes some point and reaches some critical level. The overall behavior then undergoes a *phase transition*. In our artificial stock market at low rates of investors' exploring new forecasts, the market behavior collapses to a

---

[18] Cf. Soros's (1987) observation that "stock markets are places where different propositions are tested."

[19] The literature on networks is large: see for example Albert *et al.* (2002), Allen and Gale (2000), May *et al.* (2008), Newman *et al.* (2006). Networks can be mutually stabilizing (as with banks providing insurance to other banks), but they can also be mutually destabilizing (as when losses cascade across financial institutions). And the topology of the network matters to how swiftly events propagate and to whether connectedness enhances stability or not (Scheffer *et al.*, 2012).

rational expectations equilibrium (agents make identical forecasts that produce price changes that on average validate those forecasts): simple behavior reigns. But if our investors explore at a faster, more realistic rate, the market develops a "rich psychology" of differing forecasting beliefs and starts to display temporal phenomena: complex behavior reigns. If we tune the rate of exploration still higher, individual behavior cannot adjust usefully to the rapidly changing behaviors of others, and chaotic behavior reigns. Other studies (e.g. Hommes, 2009; Kopel, 2009; LeBaron *et al.*, 1999) have found similar regime transitions from equilibrium to complexity to chaos, or from equilibrium to complexity to multiple equilibria (Galla and Farmer, 2012). Such transitions I believe will be general in nonequilibrium models.

We can now see how such phenomena—or order, or structures, if you like—connect with complexity. Complexity, as I said, is the study of the consequences of interactions; it studies patterns, or structures, or phenomena, that emerge from interactions among elements—particles, or cells, or dipoles, or agents, or firms. It's obvious that interaction takes place in our network example, but in our stock market, interaction is more subtle. If one of our investors buys or sells, this changes the price, perhaps slightly, and the others may react to this change. In all three examples, changes can propagate through the system.

Complexity studies how such changes play out. Or, to put it another way, complexity studies the propagation of change through interconnected behavior. When a bank comes under stress, it may pass this change to its connected neighbors, which may pass it to their neighbors, which may pass it to *theirs*. An event occurring at one node will cause a cascade of events: often this cascade or avalanche propagates to affect only one or two further elements, occasionally it affects more, and more rarely it affects many. The mathematical theory of this—which is very much part of complexity theory—shows that propagations of events causing further events show characteristic properties such as power laws (caused by many and frequent small propagations, few and infrequent large ones), heavy tailed probability distributions (lengthy propagations though rare appear more frequently than normal distributions would predict)[20], and long correlations (events can and do propagate for long distances and times). Such features occur in all systems—physical, chemical, biological, geological—in which events propagate, so it is not surprising that they occur in our economic examples where propagation is important.[21] They also show up tellingly in actual economic data (Brock *et al.* 1992; LeBaron *et al.*, 1999).

And we can see something else. If the degree of interaction in such a system is changed from outside (the probability of events causing further events is increased, say, or more linkages are added), the system will go from few if any consequences to many, and from that to undying-out consequences. It will go through a phase change. All these properties are hallmarks of complexity.

We can now say why nonequilibrium connects with complexity. Nonequilibrium in the economy forces us to study the propagation of the changes it causes; and complexity is very much the study of such propagations. It follows that this economics properly lies within the purview of complexity.[22]

---

[20] Their probabilities are proportional to $\exp(-|\text{propagation-length}|)$ rather than to $\exp(-\text{propagation-length}^2)$ of large normal deviations.

[21] The reason these properties do not appear in standard economics is because it assumes that agents react to a *given* equilibrium price, not to one that fluctuates due to other agents' behaviors; so random changes individual agents make are independent and can be added together. They therefore result in normal distributions.

[22] Hence this form of economics is properly called complexity economics.

One further comment. The phenomena I've illustrated appear and disappear very much in distinct historical time or space, so we will not see them if we insist on equilibrium. And they are localized: they appear in one part of the network or the stock market, possibly to diffuse from there. They operate typically at all scales—network events can involve just a few individual nodes or they can be felt right across the economy. But usually they take place in between the micro and macro, so we can rightly call them *meso-phenomena*.[23] They are properties of the *meso-economy*.

It could still be objected that such phenomena make little difference. The standard equilibrium solution after all lies beneath and still has first-order validity. This is certainly true with our stock market model; no stock will stay at 100 times earnings for long.[24] But— and this is an important "but"—the interesting things in markets happen because of their temporal phenomena, they happen within departures from equilibrium. That, after all, is where the money is made. We could similarly say that in an ocean under the undeniable force of gravity an approximately equilibrium sea level has first-order validity. And this is certainly true. But, as with markets, in the ocean the interesting things happen not at the equilibrium sea level which is seldom realized, they happen on the surface where ever-present disturbances cause further disturbances. That, after all, is where the boats are.

I have used three fairly well known phenomena in this section as illustration. Other phenomena have been noticed and no doubt others remain to be discovered. Exactly what these might be, what their characteristics are, and how they might interact are important questions for future work. But most important, our argument tells us that we need to pay attention to a new level in the economy, the meso-level, where events can trigger other events at all scales. The economy has a middle or meso layer, and in this layer phenomena arrive, last for a while, and dissipate.

Positive Feedbacks

I want to point out a further thing about the mechanisms we've been looking at. They arise from self-reinforcing behavior in the interactions. Agents buy into a stock, or disturb a market slightly, or propagate some change, and this causes further buying in, or further disturbance, or further propagation of change. Or as we saw earlier, agents show uncertainty in choice and this causes further uncertainty, or bring on some novel technology and this calls for further novel technologies. Such positive feedbacks disturb the status quo, they cause nonequilibrium. And they cause structures to appear. A small backup in traffic causes further backup and a structure forms, in this case a traffic jam. This is where the Brownian motion I alluded to comes in; it brings perturbations around which small movements nucleate; positive feedback magnifies them and they "lock in," in time eventually to dissipate.

Positive feedbacks in fact are very much a defining property of complex systems—or I should say more accurately, the presence of positive and negative feedbacks acting together is. If a system contains only negative feedbacks (in economics, diminishing returns) it quickly converges to equilibrium and shows "dead" behavior. If it contains only positive feedbacks, it runs away and shows explosive behavior. With a mixture of both it shows "interesting" or "complex" behavior. With positive feedback interactions add to each other and cause structure, in time to be offset by negative forces and dissipate. Structures then

---

[23] For earlier uses of "meso" in economics, see Dopfer (2007) and Elsner and Heinrich (2009).

[24] But it is not true in general: many economic situations do not have forces leading to any equilibrium attractor.

come and go, some stay to be further built on and some lead to further structures. The system is "alive."

These observations add to the earlier literature in economics on positive feedbacks or increasing returns. Here, if a firm (or product or technology or geographical region) gets ahead, possibly by small chance events, given increasing returns it will gain further advantage and get further ahead; it may then subsequently go on to dominate the outcome (Arthur, 1989, 1994). If N firms compete there are N possible outcomes, but N need not be small. In the late 1800s, typewriter keyboard layouts "competed" for use, and only the one we use today became a standard. But a simple calculation shows there were more than 10^{54} outcomes possible, and this is a large number by any measure.

The process that increasing returns bring into being is by now well known. What I would add is that positive feedbacks are present more widely in the economy than we previously thought: they show up not just with firms or products, but in small mechanisms and large, in decision behavior, market behavior, financial behavior, and network dynamics. They act at all scales to destabilize the economy, even the macro-scale (Keynes' theory can be seen as positive feedbacks temporarily locking in one of two possible states: full employment and unemployment). And they lead to a set of characteristic properties: multiple attractors, unpredictability, lock-in, possible inefficiencies, and path-dependence. Their counterparts in physics are multiple metastable states, unpredictability, phase- or mode-locking, high-energy ground states, and non-ergodicity. Once again these are properties we associate with formal complexity.

The Economy in Formation

I want to turn now to a very different topic, one that has builds on our earlier issue of disruption by technologies. Until now, we have seen *given* elements that comprise the economy reacting to the patterns they create and forming ever different patterns. But this still doesn't quite capture one fundamental feature of the economy. The economy continually creates and re-creates itself, and it does this by creating *novel* elements—often novel technologies and institutions—which produce novel structures as it evolves. How exactly does this happen? How does the economy form itself and change structurally? Schumpeter (1908) called this question "the most important of all the phenomena we seek to explain." Complexity should be able to help here; it is very much about the creation and re-creation of structure.

Let us begin by observing that if we want to look at how the economy constructs itself and changes, we need to look at technology and how *it* constructs itself and changes over time. Technology isn't the only agent of change in the economy but it is by far the main one (Solow, 1957). The standard story of economic change equates technologies with production functions and sees the economy as a container for these. As new industrial technologies enter, production functions change, output increases and labor or other resources are released; this provides further wealth that can be invested in further technologies. The economy shifts smoothly from one equilibrium to another and endogenously grows. This is fine and it fits well with equilibrium economics. But it puts the main driver, technology, in the background, with prices and quantities in the foreground. And it sees technologies as formless; they just somehow arrive, singly and randomly, with no structure to how they build out or how they change the economy in character over time.

A complexity view would put technologies in the foreground, and prices and quantities in the back.[25] It would recognize that there is considerable structure to how technologies arise and enter the economy (Arthur, 2009). In doing this it would focus directly on the collection of technologies present at any time, and ask how this collection evolves: how its members come into being, how they create and re-create a mutually supporting set, and how this alters the economy structurally over time.

To start, we can define individual technologies as means to human purposes. These would include industrial processes, machinery, medical procedures, algorithms, and business processes. And they would also include organizations, laws and institutions—these too are means to human purposes. The significant thing about technologies is that they are constructed, put together, combined—always—from parts, assemblies, sub-assemblies. These latter are also means to purposes, so novel technologies form by combination from existing technologies.[26] The laser printer was constructed from the existing laser, digital processor, and xerography (the processor directs a highly-focused laser beam to "paint" an image on a copier drum). We now have a system where novel elements (technologies) constantly form from existing elements, whose existence may call forth yet further elements.

Next let us define the economy as the set of arrangements and activities by which a society fulfills its needs. These arrangements of course are the economy's technologies. This is not a familiar way to look at the economy, but it fits well with the classical economists' view of the economy as proceeding from its instruments of production. The economy we can then say *emerges* from its arrangements, its technologies: it is an *expression* of its technologies. Seen this way, the economy immediately becomes an ecology of its means of production (its technologies), one where the technologies in use need to be mutually supporting and economically consistent.

We can add one more observation. Technologies come into being only if there exists a "demand" for them. Most of this demand comes from the needs of technologies themselves. The automobile "demands" or calls forth the further technologies of oil exploration, oil drilling, oil refining, mass manufacture, gasoline distribution, and car maintenance. At any time then there is an open web of opportunities inviting further technologies and arrangements.

We now have the basic setup. To put it in motion we can ask how the collection builds out. The steps involved yield the following algorithm for the formation of the economy.

1. A novel technology appears. It is created from particular existing ones, and enters the active collection as a novel element.

2. The novel element becomes available to replace existing technologies and components in existing technologies.

3. The novel element sets up further "needs" or opportunity niches for supporting technologies and organizational arrangements.

4. If old displaced technologies fade from the collective, their ancillary needs are dropped. The opportunity niches they provide disappear with them, and the elements that in turn fill these may become inactive.

---

[25] For other complexity approaches to formation see Hildago and Hausmann (2009), and Lane *et al.* (2009). On structural change see North (1981).

[26] Schumpeter (1912) cites combination as the key driving force of formation (or "development" as he called it).

5. The novel element becomes available as a potential component in further technologies—further elements.

6. The economy—the pattern of goods and services produced and consumed—readjusts to these steps. Costs and prices (and therefore incentives for novel technologies) change accordingly.

Thus the railway locomotive was constructed from the already-existing steam engine, boiler, cranks, and iron wheels. It entered the collective around 1829 (step 1); replaced existing horse-drawn trains (step 2); set up needs for the fabrication of iron rails and the organization of railways (step 3); caused the canal and horse-drayage industries to wither (step 4); became a key component in the transportation of goods (step 5); and in time caused prices and incentives across the economy to change (step 6). Such events may operate in parallel: new opportunities for example appear almost as soon as a new technology appears.

If you play the algorithm out in your mind you see something interesting. It can set in motion a sequence of happenings that never end, because each of the events may trigger a cascade of further events. A novel technology may cause further technologies to be added, by steps 3 and 5; further replacements of old technologies, by step 4; and further readjustments, by step 6. And these new technologies in turn can cause yet further opportunities, further technologies, and further replacements. The algorithm may be simple, but once set in motion it engenders rich, patterned, endlessly novel behavior.

So far this depicts the basic mechanism of formation of the economy. But there is a second layer of mechanism that adds further structure. New technologies often enter in groups (Perez, 2002; Arthur, 2009): over decades, families of technologies, the steam-driven ones, electrical ones, chemical ones, digital ones, enter. These are based on a given key technology, the steam engine say, or on families of related phenomena—chemical, electrical, genetic—that are harnessed and become available. And they build haltingly from one or two early central technologies then fill in the needed sub-technologies. These bodies of technology are not adopted within the economy, rather they are *encountered* by industries, combining with business processes that already exist and causing new activities, new incentives, new available processes, and little irruptions in the shape of little firms, a few of which go on to become large firms.

The economy—the set of arrangements and activities that satisfy our needs—builds out as a result of all this. Indeed the economy is the result of all this.

I have given only the bare bones of the processes by which the economy re-forms itself, and each mechanism has sub-mechanisms omitted here (see Arthur, 2009). But notice the overall theme: A few simple properties of technology yield a system of changing elements (technologies), each new element created from previous elements, each causing replacements, and all bringing on an ever-changing set of demands for further elements, the whole channeled and structured by the properties and possibilities of the dominant families of phenomena recently captured.

This overall process is a self-creating one. Novel technologies form from existing technologies, so the collective of technology is self-producing or *autopoietic*. So too is the economy. It forms from its technologies and mediates the creation of further technologies and thereby its own further formation. Here again we are very much in complexity territory.

We can now see how the economy changes structurally. As novel physical technologies enter, novel forms of organization and novel institutions are called for and come into place,

and these in turn call forth further new technologies—further methods, organizations, and institutions. Structure emerges. On a longer time scale, the large bodies of technology define a thematic way by which operations in the economy are carried out. So we have the steam era, the railroad era, the digital era. They also pose characteristic or thematic challenges that call forth novel solutions; the economy changes structurally. The steam engine and early textile machinery made possible the Victorian mill-based economy, and its excesses called forth new arrangements: laws covering child safety, regulations ameliorating working conditions, and labor unions in modern form.[27] As the economy changes then, its organizations and institutions change, and these call forth yet further arrangements—further technologies—and further changes. The economy transforms structurally. We can isolate the mechanisms by which the economy renews itself, but we can't predict the exact ways these play out. The overall process (or computation, if you will) is far from determinate. And it is par excellence one of nonequilibrium.

Notice that the theory I have outlined is algorithmic: it is expressed as a set of processes triggered by other processes, not as a set of equations. The reader may again ask how this can be theory? Consider a parallel with biology. Even now, 150 years after Darwin's *Origin*, no one has succeeded in reducing to an equation-based system the process by which novel species are created, form ecologies, and bring into being whole eras dominated by characteristic species. The reason is that the evolutionary process is based on mechanisms that work in steps and trigger each other, and it continually defines new categories—new species. Equations do well with changes in number or quantities within *given* categories, but poorly with the appearance of new categories themselves. Yet we must admit that evolution's central mechanisms are deeply understood and form a coherent group of general propositions that match real world observations, so these understandings indeed constitute theory.[28] Biology then is theoretical but not mathematical; it is process-based, not quantity-based. In a word it is *procedural*. By this token, a detailed economic theory of formation and change would also be procedural.[29] It would seek to understand deeply the mechanisms that drive formation in the economy and not necessarily seek to reduce these to equations. The procedural theory I have outlined doesn't negate the standard one, but it does give an alternative that puts the emphasis squarely on the driver of change itself—on technology.

How can we study all this more deeply? The base processes are algorithmic, so certainly we can construct computer-based models of their key mechanisms.[30] Studies here are still at

---

[27] Political economist William Tabb (1999), expresses structural change this way: "Technological revolutions and political upheavals condition economic possibilities, which then become the givens for sustained periods of seeming stability in which regulatory regimes designed for the conditions of the social structure of accumulation of the era lend a semblance of orderly progress. These institutional forms, appropriate to one stage of development, become a drag on the development of new forces and emergent relations of production. The vitality of market forces create in their wake social problems which, when they become severe enough need to be addressed through spirited struggle out of which new rules, regulations, and institutions form."

[28] Similar observations can be made about the theories of embryological development, of biochemical pathways, of molecular genetics, and of cell biology. The process of mitosis (cell division) has no mathematics, but does have a series of well-understood, if complicated, phases or steps.

[29] The reader might be tempted to translate this back into familiar terms such as capital, labor, growth, etc. That might be possible, but I prefer to see this as a different way to "image" or understand change in the economy, much as MRI scanning images organs differently than conventional x-ray scanning.

[30] In 2006 Wolfgang Polak and I modeled a creation process successfully on the computer by which increasingly complicated technologies (digital logic circuits) emerged from initially simple ones via random combination of earlier combinations (circuits).

their beginning. The overall view we end up with is one of creative formation: of new elements forming from existing elements, new structure forming from existing structure, formation itself proceeding from earlier formation. This is very much a complexity view.

Discussion

It should be clear by now that we have a different framework for thinking about the economy, one that emphasizes not the physics of goods and services, but processes of change and creation. Yet, as the reader may have surmised, this new view is not entirely new within economic thought. It links with earlier thinking in a way I want to comment on now.

There are two great problems in economics. One is *allocation* within the economy: how quantities of goods and services and their prices are determined within and across markets. This is represented by the great theories of general equilibrium, international trade, and game-theoretic analysis. The other is *formation* within the economy: how an economy emerges in the first place, and grows and changes structurally over time. This is represented by ideas about innovation, economic development, structural change, and the role of history, institutions, and governance in the economy. The allocation problem is well understood and highly mathematized, the formation one less well understood and barely mathematized.[31]

How did this come about? Until about 1870 both problems were of equal importance to the great theorists in economics. Smith, Mill, and Marx all contributed to making a rational science out of allocation, yet they equally contributed to questions of formation, governance, and history. Then in Victorian times came the great marginalist and general-equilibrium revolution that rendered the problem of allocation into algebra and calculus (given strict assumptions of rationality and equilibrium). But the problem of formation could not be so rendered. By its nature it couldn't be restricted to either stasis or rationality, and so the mathematization of economics—what came in the twentieth century to be taken as "theory"—passed it by. Formation was still studied by Marshall, Veblen, Schumpeter, Hayek, and Shackle, and by the many institutionalists and historians that followed. But the thinking was largely history-specific, particular, case-based, and intuitive—in a word, literary—and therefore held to be beyond the reach of generalizable reasoning, so in time what had come to be called political economy became pushed to the side, acknowledged as practical and useful but not always respected.

It is now clear to economists that the mathematical analysis of allocation far from covers all of economics and operates poorly with questions of formation, exploration, adaptation, and qualitative change (Tabb, 1999). Complexity economics by contrast is very much about these questions of creation and the formation of structure, and it studies the mechanisms by which these operate. So here complexity meets up with and revives the grand tradition of political economy, and the two—much to my delight—have a lot to say to each other. Complexity economics allows us to explore the world of formation theoretically and systematically; political economy allows us to explore it intuitively and empirically. The new approach will help provide a theoretical backbone for political economy. It will not and should not displace case-based historical analysis, but will deepen and develop this venerable branch of thinking. And political economy will deepen and develop complexity economics.

---

[31] See Tabb (1999) for an excellent discussion of these two branches of economics. Also Bronk (2009).

One of the main strengths of political economy is its sense of history, of historical time—time that makes a real, irreversible difference, and that continually creates new structures. By contrast neoclassical economics handles time poorly (Smolin, 2009, 2012). At equilibrium an outcome simply persists and so time largely disappears; or in dynamic models it becomes a parameter that can be slid back and forth reversibly to denote the current outcome (Harris, 2004). This has made many economic thinkers uncomfortable (Robinson, 1980). In 1973 Joan Robinson said famously, "Once we admit that an economy exists in time, that history goes one way, from the irrevocable past into the unknown future, the conception of equilibrium … becomes untenable. The whole of traditional economics needs to be thought out afresh."

Certainly, in rethinking this issue of time, complexity economics accords with political economy. In the "computation" that is the economy, large and small probabilistic events at particular non-repeatable moments determine the attractors fallen into, the temporal structures that form and die away, the technologies that are brought to life, the economic structures and institutions that result from these, the technologies and structures that in turn build upon *these*; indeed the future shape of the economy—the future path taken. The economy at all levels and at all times is path dependent. History again becomes important. And time reappears.

A natural question is whether this new approach has policy implications. Certainly, complexity teaches us that markets left to themselves possess a tendency to bubbles and crashes, induce a multiplicity of local attractor states, propagate events through financial networks, and generate a sequence of technological solutions and challenges, and this opens a role for policies of regulating excess, nudging towards favored outcomes, and judiciously fostering conditions for innovation. Colander and Kupers (2012) express this succinctly as getting meta conditions right.

This is certainly valid. But I believe we can make a stronger statement. The failures of economics in the practical world are largely due to seeing the economy in equilibrium. If we look at the economic crises of the last 25 years—the debacle that followed the freeing of markets in Russia in 1990, the extensive gaming of California's energy market after the lifting of regulations in 2000, the collapse of Iceland's banks in 2008, the ongoing Euro crisis, the Wall Street meltdown of 2008—all these were caused in no small part by the exploitation of the system by a few well-positioned players, or by markets that careened out of control (Arthur, 2010a). Equilibrium thinking cannot "see" such exploitation in advance for a subtle reason: by definition, equilibrium is a condition where no agent has any incentive to diverge from its present behavior, therefore exploitive behavior cannot happen. And it cannot see extreme market behavior easily either: divergences are quickly corrected by countervailing forces. By its base assumptions, equilibrium economics is not primed to look for exploitation of parts of the economy or for system breakdowns.

Complexity economics, by contrast, teaches us that the economy is permanently open to response and that every part of it is open to new behavior—to being exploited for gain, or to abrupt changes in structure. A complexity outlook would recommend putting carefully thought out controls in place, much as authorities put sensible building codes in place in seismic regions. But just as important, it would bring a shift in attitude in the direction of realism. The economy does not consist of a set of behaviors that have no motivation to change and collectively cause optimality; the economy is a web of incentives that always induce further behavior, invite further strategies, provide collectively "reasonable" outcomes along the way, and ever cause the system to change.

Conclusion

Complexity economics is neither an add-on to standard economics (see Fontana, 2010), nor does it consist of adding agent-based behavior to standard models. It is a different way of thinking about the economy. It sees the economy not as a system in equilibrium but as one in motion, perpetually "computing" itself—perpetually constructing itself anew. Where equilibrium economics emphasizes order, determinacy, deduction, and stasis, this new framework emphasizes contingency, indeterminacy, sense-making, and openness to change. There is another way to say this. Until now, economics has been a noun-based rather than verb-based science. It has pictured changes over time in the economy function as changes in levels of fixed noun-entities—employment, production, consumption, prices. Now it is shifting toward seeing these changes as a series of verb-actions—forecast, respond, innovate, replace—that cause further actions.

This shift reveals an important middle layer in the economy, the *meso-layer*. And it redefines what constitutes a solution in economics. A solution is no longer necessarily a set of mathematical conditions but a pattern, a set of emergent phenomena, a set of changes that may induce further changes, a set of existing entities creating novel entities. Theory in turn becomes not the discovery of theorems of undying generality, but the deep understanding of mechanisms that create these patterns and propagations of change.

This shift in economics is very much part of a larger shift in science itself. All the sciences are becoming more procedural, more algorithmic, more Turingesque; and less equation-based, less continuous, less Newtonian, than before. This is due both to the rise of biology as a rigorous science and to the rise of computation and computer science. Even mathematics is shifting in this direction. Gregory Chaitin (2012) speaks of a mathematics that is shifting away from continuous formulations, differential equations, and static outcomes, to one based on discrete formulations, combinatorial reasoning, and algorithmic thinking. "The computer," he says, "is not just a tremendously useful technology, it is a revolutionary new kind of mathematics with profound philosophical consequences. It reveals a new world." Science and mathematics are shedding their certainties and embracing openness and procedural thinking, and there is no reason to expect that economics will differ in this regard.

Complexity economics is not a special case of neoclassical economics. On the contrary, equilibrium economics is a special case of nonequilibrium and hence complexity economics. Complexity economics, we can say, is economics done in a more general way. Equilibrium of course will remain a useful first-order approximation, useful for situations in economics that are well-defined, rationalizable, and reasonably static, but it can no longer claim to be the center of economics. Moving steadily to the center[32] is an economics that can handle interactions more generally, that can recognize nonequilibrium phenomena, that can deal with novelty, formation and change.

Complexity economics is still in its early days and many economists are pushing its boundaries outward. It shows us an economy perpetually inventing itself, perpetually creating possibilities for exploitation, perpetually open to response. An economy that is not dead, static, timeless, and perfect, but one that is alive, ever-changing, organic, and full of messy vitality.

---

[32] See Holt *et al.* (2010); Davis (2007).

References

ACM Ubiquity Symposium: *What is Computation?* 2010.

Angeletos, G-M, and J. La'O, "Decentralization, Communication, and the Origins of Fluctuations," NBER Working Paper 17060, NBER, Cambridge, Ma., 2011.

Albert, R., H. Jeong, and A-L. Barabasi, "Attack and Error Tolerance of Complex Networks," *Nature,* 406, 379-82, 2000.

Allen, F. and D Gale, "Financial Contagion," *Journal of Political Economy*, 108, 1. 1-33, 2000.

Arrow, K., P. Anderson and D. Pines, *The Economy as an Evolving Complex System*, Addison-Wesley, Reading, Ma., 1988.

Arthur, W.B., "Competing Technologies, Increasing Returns, and Lock-In by Historical Events," *Economic Journal*, 99, 116-31, 1989.

Arthur, W.B., "Bounded Rationality and Inductive Behavior (the El Farol problem)," *American Economic Review Papers and Proceedings*, 84, 406-411, 1994.

Arthur, W.B., *Increasing Returns and Path Dependence in the Economy*, U Mich Press, Ann Arbor, Mi., 1994.

Arthur, W.B., S. Durlauf, D. Lane, eds. *The Economy as an Evolving Complex System II*, Addison-Wesley, Reading, Ma., 1997.

Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer, P. Tayler, "Asset Pricing under Endogenous Expectations in an Artificial Stock Market," in *The Economy as an Evolving Complex System II, (op. cit.*), 1997.

Arthur, W.B., 'Complexity and the Economy,' *Science*, 284, 107-109, 1999.

Arthur, W.B., "Out-of-equilibrium Economics and Agent-based Modeling," in L. Tesfatsion and K. Judd, (*op. cit.* below), 2006.

Arthur, W.B. and W. Polak, "The Evolution of Technology in a Simple Computer Model," *Complexity*, 11, 5, 2006.

Arthur, W.B., *The Nature of Technology: What it Is and How it Evolves,* Free Press, New York, 2009.

Arthur, W.B., "Exploitive Behavior in Policy Systems," Mss., IBM Almaden, 2010a.

Arthur, W.B., "Complexity, the Santa Fe Approach, and Nonequilibrium Economics," *Hist. Econ. Ideas,* 18, 2, 149-66, 2010b.

Axtell, R., "What Economic Agents Do: How Cognition and Interaction Lead to Emergence and Complexity," *Rev. Austrian Econ.*, 20, 105–22, 2007.

Bailey, J., *Emerge: The Data-Rich Mathematical Infinitesimals of Life*, MAPematics, 2010.

Beinhocker, E., *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics,* Harvard Business School Press, Cambridge, Ma., 2006.

Beinhocker, E., "Evolution as Computation: Integrating Self-Organization with Generalized Darwinism," *J. Inst. Econ.*, 7, 3, 393-423, 2011.

Blaug, M., "The Formalist Revolution of the 1950s," *J. History of Economic Thought*, 25, 2, 145-56, 2003.

Blume, L. and S. Durlauf, *The Economy as an Evolving Complex System III*, Oxford U Press, NY, 2006.

Brock, W.A., J. Lakonishok, and B. LeBaron, "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns," *Journal of Finance,* 47, 1731-64, 1992.

Bronk, R., *The Romantic Economist: Imagination in Economics*, Cambridge U. Press, Cambridge, UK, 2009.

Bronk, R., "Epistemological Difficulties with Neoclassical Economics," *Rev. Austrian Economics*, forthcoming 2013.

Cassidy, J., *How Markets Fail: the Logic of Economic Calamities*, Farrar, Straus and Giroux, NY, 2009.

Chaitin, G., *Meta Math! The Quest for Omega,* Vintage Books, NY, 2006.

Chaitin, G., *Proving Darwin: Making Biology Mathematical*, Pantheon Books, NY, 2012.

Colander, D, and R. Kupers. *Laissez-Faire Activism: The Complexity Frame for Policy*, Princeton, 2012.

Colander, D. (ed.), *The Complexity Vision and the Teaching of Economics*, E. Elgar, Cheltenham, UK, 2000.

Colander, D., M. Goldberg, A. Haas, K. Juselius, T. Lux, H. Föllmer, A. Kirman, B. Sloth, "The Financial Crisis and the Systemic Failure of the Economics Profession," *Critical Rev.,* 21, 2, 2009.

Davis, J., "The Turn in Recent Economics and Return of Orthodoxy," *Cambridge J. Econ.,* 32, 349–66, 2008.

Dopfer, K., "The Pillars of Schumpeter's Economics: Micro, Meso, Macro," in Hanusch and Pyka, eds., *Elgar Companion to Neo-Schumpeterian Economics.* E. Elgar, Cheltenham, UK., 2007.

Elsner W. and T. Heinrich, "A Simple Theory of 'Meso:' on the Co-evolution of Institutions and Platform Size," in *Journal of Socio-Economics*, 38, 843-58, 2009.

Epstein J., *Generative Social Science,* Princeton Univ. Press, Princeton, NJ, 2006a.

Epstein, J., "Remarks on the Foundations of Agent-based Generative Social Science," in Tesfatsion and Judd, eds., *op. cit.* (below)*,* 2006b.

Farmer, J. D. and J. Geanakoplos, "The Virtues and Vices of Equilibrium and the Future of Financial Economics," *Complexity*, 14, 8, 11-38, 2008.

Farmer, J.D., "Economics Needs to Treat the Economy as a Complex System," Mss., 2012.

Fontana, M., "Can Neoclassical Economics handle Complexity? The Fallacy of the Oil Spot Dynamic," *J. Economic Behavior & Organization*, 76, 584-596, 2010.

Fontana, M. "The Santa Fe Perspective on Economics," *Hist. Econ. Ideas*, 18, 2, 167-96, 2010.

Galla, T. and J.D. Farmer, "Complex Dynamics in Learning Complicated Games," Mss., 2012.

Haldane, A.G., "Rethinking the Financial Network," Speech at Financial Student Assoc., Amsterdam. Bank of England, 2009.

Harris, D.J., "Joan Robinson on 'History versus Equilibrium,'" J. Robinson Centennial Conference, Burlington, Vt., 2003.

Hildago C.A. and R. Hausmann, "The Building Blocks of Economic Complexity," *Proc. Nat. Acad. Sci.,* 106, 26, 10570-75, 2009.

Holland, J., K. Holyoak, R. Nisbett and P. Thagard, *Induction*, MIT Press, Cambridge, Ma., 1986.

Holt R., J. Rosser, and D. Colander, "The Complexity Era in Economics," Middlebury College, 2010.

Hommes, C.H. "Bounded Rationality and Learning in Complex Markets," in Rosser, J.B. (*op. cit.* below), 2009.

Judd, K., "Computationally Intensive Analyses in Economics," in Tesfatsion and Judd, eds., (*op. cit.* below), 2006.

Keynes, J.M., "The General Theory of Employment," *Quarterly Journal of Economics*, 51, 209-33. 1937.

Kirman, A., *Complex Economics*, Routledge, NY, 2011.

Kirman, A., "The Economic Crisis is a Crisis for Economic Theory," *CESifo Economic Studies*, 56, 4, 498-535, 2010.

Kopel, M., "Oligopoly Dynamics," in Rosser, J.B., (*op. cit.* below), 2009.

Koppl, R., W. Luther, "BRACE for a new Interventionist Economics," Fairleigh Dickinson U., 2010.

Krugman, P., "How did Economists Get it so Wrong?" *NY Times,* Sep 6, 2009.

Lane, D., D. Pumain, S. van der Leeuw, and G. West (eds.), *Complexity Perspectives in Innovation and Social Change,* Springer, Berlin, 2009.

LeBaron, B., W.B. Arthur, and R. Palmer, "Time Series Properties of an Artificial Stock Market, *Journal of Econ. Dynamics and Control*, 23, 1487-1516, 1999.

Lindgren, K., 'Evolutionary Phenomena in Simple Dynamics,' in C. Langton, C. Taylor, J.D. Farmer and S. Rasmussen, eds, *Artificial Life II*, Addison-Wesley, Reading Ma., 1991.

Louça, F., "Bounded Heresies. Early Intuitions of Complexity in Economics." *Hist. Econ. Ideas,* 18, 2, 77-113, 2010.

May, R., S. Levin, G. Sugihara, "Complex Systems: Ecology for Bankers," *Nature*, 451, 893-5, 2008.

Miller J., and S. Page, *Complex Adaptive Systems: An Introduction to Computational Models of Social Life,* Princeton U. Press, Princeton, NJ., 2007.

Mirowski, P., *Machine Dreams: Economics Becomes a Cyborg Science*, Cambridge U. Press, UK., 2002.

Newman, M., A-L. Barabasi, D. Watts (eds.). *The Structure and Dynamics of Networks,* Princeton, 2006.

North, D., *Structure and Change in Economic Theory*, Norton, NY, 1981.

Palmer, R.G., W.B. Arthur, J. Holland, B. LeBaron, P. Tayler, "Artificial Economic Life: A Simple Model of a Stock market," *Physica D,* 75, 264-74, 1994.

Perez, C., *Technological Revolutions and Financial Capital*, E. Elgar, Cheltenham, UK, 2002.

Reisman, D., *Schumpeter's Market: Enterprise and Evolution*, E. Elgar, Cheltenham, UK, 2004.

Robinson, J., "Time in Economic Theory," *Kyklos,* 33, 2. 219-29, 1980.

Robinson, J., "What has Become of the Keynesian Revolution?" in J. Robinson, ed. *After Keynes*, Basil Blackwell, Oxford, 1973.

Robertson, D.S. *Phase Change: the Computer Revolution in Science and Mathematics*, Oxford, NY, 2003.

Rosser, J.B., "On the Complexities of Complex Economic Dynamics," *J. Econ Perspectives,* 13, 4, 1999.

Rosser, J.B., ed. *Handbook of Research on Complexity,* E. Elgar. Cheltenham, UK, 2009.

Samuelson, P.A., *Foundations of Economic Analysis*, Harvard, Cambridge, MA., 1983 (originally 1947).

Sargent T. J., *Bounded Rationality in Macroeconomics*, Clarendon Press, Oxford, 1993.

Scheffer, M., S. Carpenter, T. Lenton, J. Bascompte, W. Brock, V. Dakos, J. van de Koppel, I. van de Leemput, S. Levin, E. van Nes, M. Pascual, J. Vandermeer, "Anticipating Critical Transitions," *Science,* 338, 19 Oct. 2012.

Schumpeter, J.A., *Das Wesen und der Hauptinhalt der theoretischen Nationalökonomie*. Dunker & Humbolt, Leipzig, 1908.

Schumpeter, J.A., *The Theory of Economic Development*, (1912), Oxford Univ. Press, London, 1961.

Schumpeter, J.A., *History of Economic Analysis,* Allen and Unwin, London, 1954.

Shackle, G.L.S., *Uncertainty in Economics*, Cambridge U. Press, Cambridge, UK, 1955.

Shackle, G.L.S., *Epistemics and Economics,* Transaction Publishers, 1992.

Simpson, D., *Rethinking Economic Behaviour*, St. Martin's Press, NY, 2002.

Smolin, L., "Time and Symmetry in Models of Economic Markets," Mss., Feb 2009.

Smolin, L., *Time Reborn*. Houghton, Mifflin, Harcourt, NY, 2013.

Solow, R., "Technical Change and the Aggregate Production Function," *Rev. Economics and Statistics,* 39, 312-20, 1957.

Soros, G., *The Alchemy of Finance*, Simon & Schuster, NY, 1987.

Tabb, W., *Reconstructing Political Economy*, Routledge, NY, 1999.

Tesfatsion, L. and K.L. Judd, eds., *Handbook of Computational Economics: Vol. 2. Agent-based Computational Economics*, North-Holland Elsevier, NY, 2006.

Tesfatsion, L., "Agent-Based Computational Economics: A Constructive Approach to Economic Theory," in Tesfatsion and Judd, *op. cit.*, 2006.

Turing, A.M., "On Computable Numbers, with an Application to the Entscheidungsproblem," *Proc. London Math. Society*, Series 2, 42, 1936.

Waldrop, M., *Complexity,* Simon & Schuster, NY, 1992.

Watts, D., "A Simple Model of Global Cascades on Random Networks," *PNAS*, 9, 5766-71, 2002.

Wolfram, S., *A New Kind of Science*, Wolfram Media, Champaign, Il., 2002.