COMPLEXITY

and the ECONOMY

W. BRIAN ARTHUR

Oxford University Press

CONTENTS

PREFACE

Every so often a discipline gets thrown into a period of upheaval where its old ideas once taken for granted seem no longer so reliable, and its practitioners search for what to put in their place. Economics is in such a period now. This is partly due to the financial crisis of 2008, but the rethinking goes back to well before this. Slowly, over the last three or more decades, a feeling has grown among economists that their key assumptions of perfect rationality, equilibrium, diminishing returns, and of independent agents always facing well-defined problems are somehow not trustworthy, too restrictive, somehow forced. Now in the air are ideas of behavioral rationality, nonequilibrium, increasing returns, and of interconnected agents facing fundamental uncertainty in problems of decision-making. Economics has opened up to other approaches besides the standard neoclassical one.

I have been heavily involved in one of the new approaches, complexity economics, so I decided this would be a good time to put together several of my earlier papers and bring them out in a collected volume. This collection, on complexity and the economy, dates from the mid-1980s to the present, and it follows my earlier one on increasing returns and path dependence in the economy.[1]

None of these "new" ideas of course are really completely that new. Separately and in various forms they have been mooted by economists for years, sometimes even for a century or more. But what has been missing was the means to handle them, not just raw techniques but the mindset that would go with them, that the world is not perfect, that it isn't machine-like, and that much of it cannot be reduced to simple equations—to variations in the number or level of entities. And missing too was a coherent framework for economics based on these new ideas.

In the last few decades this has changed. The missing pieces have begun to fill in and techniques have slowly become available that can deal with the new assumptions. Among these are nonlinear dynamics, nonlinear stochastic processes, agent-based computation, and computational theory itself. The mindset too has changed. A feeling now runs across the sciences, and economics too, that the world is not a perfectly ordered system reducible in principle to mathematical equations, but is to a large extent organic and algorithmic—it proceeds by building on what is there already and it builds and changes step by step. Slowly, as a result of these occurrences, economics is

[1] *Increasing Returns and Path Dependence in the Economy*, W. B. Arthur, Univ. Mich. Press, 1994.

developing an approach based on these more realistic assumptions. It is developing a new framework for economic thought.

The collected papers in this book reflect my part in the development of this new framework. Taken together they view the economy as a system not necessarily in equilibrium, but as one where agents constantly change their actions and strategies in response to the outcome they mutually create, a system where agents are constantly creating an "ecology" of behaviors they must mutually adapt to. This viewpoint has roots of course in complexity thinking as it developed in the 1970s in groups in Brussels, Stuttgart, and Ann Arbor. And it has roots in the work of individual researchers in universities such as Stanford and MIT. But in its current economic form it grew largely from work at the Santa Fe Institute. In the late 1980s a small group of researchers at the nascent Santa Fe Institute began systematically to look at the economy as an evolving complex system. I headed that group for its first two years, and have been associated with it ever since, and in this collection of papers I want to show how these ideas developed and how the economics they led to came about.

The papers in this volume were not the outcome of some planned process. They arose haltingly and over several years, and were heavily influenced by my colleagues and by thinking in general at Santa Fe. Several appeared in well-known journals, others appeared in places more obscure. Many were written in Santa Fe, others were written at Stanford. The papers present finished thinking but not why or how that thinking came about, so it will be helpful to the reader to understand the background to them and the context in which they arose.

Most of them started with a single incident.

§

In April 1987 I was walking toward my office in Stanford when a helmeted Kenneth Arrow swung round me on his bicycle and stopped. He was putting together a group of economic theorists in September to exchange ideas with a group of physicists that his counterpart, physicist Philip Anderson, would propose. The venue would be a small institute in the Rockies just starting up. It was in Santa Fe. Would I like to come? I said yes immediately without being sure of what I was committing myself to. The idea looked promising.

The conference in Santa Fe when I got there a few months later turned out to be a more heavyweight affair than I'd imagined. Among the ten or so economists Arrow chose were Larry Summers, Tom Sargent, Jose Sheinkman, and William (Buz) Brock. Among the ten or so "physicists" Phil Anderson chose were John Holland, David Ruelle, Stuart Kauffman, and David Pines. The meeting was held in the chapel of a convent the new institute was renting and there was nothing rushed about it. A participant would talk in the morning and we would discuss, another participant would talk in the afternoon and again we would discuss. We were learning not just solutions to problems in the others' disciplines, but about what each discipline saw as a problem, and how it thought about these, and what mindset it brought to bear on these problems. Questions not normally raised within economics were raised—why do you guys cling onto perfect rationality? Why do you assume so much linearity? And questions were asked of physics too. Why is a problem "solved," say in spin glasses, when it has not settled to a steady state? Chaos theory and nonlinear dynamics were discussed in both economics and physics. Modeling of positive feedbacks and of interactions, again in both disciplines, was discussed. People would meet at night in twos and threes to talk over ideas and problems.

The meeting was exhilarating—and exhausting. Nothing had quite been solved by the end of the ten days, yet the physics side was left with a respect for the sheer complicatedness of the economy—the elements in the economy (people), unlike the ions in a lattice, could decide what to do next not just based on the current situation of themselves and other elements, but on what they thought those other elements might do given what *they* might do. And the economists were left with a feeling for modern physics, for its interactions and nonlinearities, its multiple possible end states, its lack of predictability—indeed for its complicatedness.

Word began to leak out after the conference that something interesting had happened at Santa Fe and the new institute's Science Board decided it would follow the conference up by initiating a long-term research program on the Economy as an Evolving Complex System. John Holland and I were asked to come to Santa Fe the following year to head this. I had a sabbatical coming from Stanford and accepted, John found it harder to get away from Michigan and declined. So I found myself heading up the Santa Fe Institute's first research program; it would start in August the following year, 1988.

My immediate problem of course, working from Stanford, was to put together a team of first-rate people for the new program and to decide its direction. Some people I already knew from the conference. John Holland promised to come for a couple of months, and the physicist Richard Palmer for much longer than that. Stuart Kauffman would be in residence. From my own network I was able to bring in David Lane and Yuri Ermoliev, both excellent probability theorists. Arrow and Anderson helped greatly. Where I found it hard to cajole people to join in, Arrow or Anderson, both Nobel Prize winners, could simply lift the phone and quickly get people to join us. As to direction I was less sure. Early on, the physicist Murray Gell-Mann suggested to me that we come up with a manifesto for doing economics differently. I didn't quite have the confidence for that; in fact I didn't yet know what topics we would go after. I had done quite a bit of work already on complexity and the economy, but now we had a much broader reach in what topics we might research. From the conference it was assumed that chaos theory would be central, but the idea somehow didn't appeal to me. Vaguely I thought that we should look at increasing returns problems, which I was more than familiar with, at how some of the physics methods could be transferred into economics, and at nonlinear dynamics in the economy. Also we might be able to do something interesting with computation in economics.

When the program opened finally in 1988 we discussed directions further, still groping for a way forward. I phoned Ken Arrow from Santa Fe and asked for his advice and Phil Anderson's. They got in touch with the funder of the program, John Reed of Citibank, and the word came back: Do what you want, providing it deals with the foundations of economics, and is not conventional. For me and the others on the team, this directive seemed like a dream. We had carte blanche to do what we wanted, and at Santa Fe we wouldn't have colleagues from the discipline looking at us and asking why we were doing things differently.

In fact, outside our small team the few colleagues we did have were from physics or theoretical biology. Stuart Kauffman was one, and we immediately included him in the program. There was little else in the way of researchers the new institute could offer. It was in its earliest days and was all but unknown, an experiment, a small startup in the Rockies set up to have no students, no classes, no departments, and no disciplines—no discipline, the wags said.

We had discussions, mainly in the convent's kitchen, and I remember in an early one Kauffman said, Why do you guys do everything at equilibrium? What would it be like to do economics out of equilibrium? Like all economists I had thought about that, but not seriously. In fact the question took me aback, and it did so with the other economists. I had no good answer. It fell into the category of questions such as what would physics be like if the gravitational force were suspended, something that seemed perfectly thinkable as a thought experiment, but strange. And yet Kauffman's question stuck. We retained the question but we were still looking for a direction ahead.

§

One of the directions that interested me was still half formed. It had come out of the conference the previous year. In an after-lunch talk the first day of that conference, John Holland had described his work on classifier systems, basically systems that are concatenations of condition-action rules. One rule might say that *if* the system's environment fulfills condition *A*, *then* execute action *R*. Another might say, *if* it fulfills condition *D*, execute action *T*. A third might say that *if A* is true, *and* R-being-executed is not true, *then* execute action *Z*. And so on. The actions taken would change the environment, the overall state of the system. In this way you could string such if-then rules together to get a system to "recognize" its environment and execute actions appropriately, much as an *E. coli* bacterium "recognizes" a glucose gradient in its environment and swims in an appropriate direction. Moreover, you could allow the system to start with not-so-good rules and replace these with better ones it discovered over time. The system could learn and evolve.

As Holland talked about this I found myself deeply excited, and I checked the room to see if other economists were similarly taken with these ideas. There was no evidence; in fact one of them was taking a post-lunch nap. A feeling grew in me that somehow, in some way, this was an answer and all we had to do was find the question. Somehow Holland was describing a method whereby "intelligence" or appropriate action could automatically evolve within systems. I quizzed John later about his ideas. We were sharing a house in Santa Fe for two months at that time in 1987, but in several conversations neither of us could work out what these ideas might directly have to do with economics.

I had gone back to Stanford, where I was teaching a course in economic development. It occurred to me, gradually at first, that John and I could design a primitive artificial economy that would execute on my computer, and use his learning system to generate increasing sophisticated action rules that would build on each other and thus emulate how an economy bootstraps its way up from raw simplicity to modern complication. In my mind I pictured this miniature economy with its little agents as sitting in a computer in the corner of my office. I would hit the return button to start and come back a few hours later to peer in and say, oh look, they are trading sheep fleeces for obsidian. A day later as the computation ran, I would look again and see that a currency had evolved for trading, and with it some primitive banking. Still later, joint stock companies would emerge. Later still, we would see central banking, and labor unions with workers occasionally striking, and insurance companies, and a few days later, options trading. The idea was ambitious and I told Holland about it over the phone. He was interested, but neither he nor I could see how to get it to work.

That was still the status the following summer in June 1988 when Holland and I met again in Santa Fe shortly before the program was to start. I was keen to have some form of this self-evolving economy to work with. Over lunch at a restaurant called Babe's on Canyon Road, John asked how the idea was coming. I told him I found it difficult, but I had a simpler idea that might be feasible. Instead of simulating the full development of an economy, we could simulate a stock market. The market would be completely stand-alone. It would exist on a computer and would have little agents—computerized investors that would each be individual computer programs—who would buy and sell stock, try to spot trends, and even speculate. We could start with simple agents and allow them to get smart by using John's evolving condition-action rules, and we could study the results and compare these with real markets. John liked the idea.

We began in the fall, with the program now started, to build a computer-based model of the stock market. Our "investors," we had decided, would be individual computer programs that could react and evolve within a computer that sat on my desk. That much was clear, but we had little success in reducing the market to a set of condition-action rules, despite a number of attempts. The model was too ad-hoc, I thought—it wasn't clean. Tom Sargent happened to be visiting from Stanford and he suggested that we simply use Robert Lucas's classic 1978 model of the stock market as a basis for what we were doing. This worked. It was both clean and doable. Lucas's model of course was mathematical; it was expressed in equations. For ease of analysis, his investors had been identical; they responded to market signals all in the same way and on average correctly, and Lucas had managed to show mathematically how a stock's price over time would vary with its recent sequence of earnings.

Our investors, by contrast, would potentially differ in their ideas of the market and they would have to learn what worked in the market and what didn't. We could use John's methods to do this. The artificial investors would develop their own condition/forecast rules (e.g., if prices have risen in the last 3 periods *and* volume is down more than 10%, *then* forecast tomorrow's price will be 1.35% higher). We would also allow our investors to have several such rules that might apply—multiple hypotheses—and at any time they would act on the one that had proved recently most accurate of these. Rules or hypotheses would of course differ from investor to investor; they would start off chosen randomly and would be jettisoned if useless or recombined to generate potential new rules if successful. Our investors might start off not very intelligently, but over time they would discover what worked and would get smarter. And of course this would change the market; they might have to keep adjusting and discovering indefinitely.

We programmed the initial version in Basic on a Macintosh with physicist Richard Palmer doing the coding. Initially our effort was to get the system to work, to get our artificial investors to bid and offer on the basis of their current understandings of the market and to get the market to clear properly, but when all this worked we saw little at first sight that was different from the standard economic outcome. But then looking more closely, we noticed the emergence of real market phenomena: small bubbles and crashes were present, as were correlations in prices and volume, and periods of high volatility followed by periods of quiescence. Our artificial market was showing real-world phenomena that standard economics with its insistence on identical agents using rational expectations could not show.

I found it exciting that we could reproduce real phenomena that the standard theory could not. We were aware at the time that we were doing something different. We were simulating a market in which individual behavior competed and evolved in an "ecology" that these behaviors mutually created. This was something that couldn't easily be done by standard equation-based methods—if forecasting rules were triggered by specific conditions and if they differed from investor to investor their implications would be too complicated to study. And it differed from other computerized rule-based models that had begun to appear from about 1986 onward. Their rules were few and were fixed—laid down in advance—and tested in competition with each other. Our rules could change, mutate, and indeed "get smart." We had a definite feeling that the computer would free us from the simplifications of standard models or standard rule-based systems. Yet we did not think of our model as computer simulation of the market. We saw it as a lab experiment where we could set up a base case and systematically make small changes to explore their consequences.

We didn't quite have a name for this sort of work—at one stage we called it element-based modeling, as opposed to equation-based modeling. About three years later, in 1991, John Holland and John Miller wrote a paper about modeling with "artificial adaptive agents."[2] Within the economics community this label morphed into "agent-based modeling" and that name stuck. We took up other problems that first year of the Economics Program. Our idea was not to try to lay out a new general method for economics, as Samuelson and others had tried to do several decades before. Rather we would take known problems, the old chestnuts of economics, and redo them from our different perspective. John Rust and Richard Palmer were looking at the double auction market this way. David Lane and I were working on information contagion, an early version of social learning, using stochastic models. I had thought that ideas of increasing returns and positive feedbacks would define the first years of the program. But they didn't. What really defined it, at least intrinsically, was John Holland's ideas of adaptation and learning. I had also thought we were going slowly and not getting much done, but at the end of our first year, in August 1989, Kenneth Arrow told us that compared with the initial years of the Cowles Foundation effort in the 1950s, our project had made faster progress and was better accepted.

I left Santa Fe and returned to Stanford in 1990 and the program passed into other hands. It continued with various directors throughout the 1990s and the early 2000s with considerable success, delving into different themes depending on the directors' interests and passing through periods of relative daring and relative orthodoxy. I returned to the Institute in 1995 and stayed with the Program for a further five years.

---

[2] J. H. Holland and J. H. Miller, "Artificial Adaptive Agents in Economic Theory," *Amer. Econ. Assoc. Papers and Proceedings,* 81, 2, 365–370, 1991.

Most of the economic papers in this volume come out of this first decade or so of SFI's economics program. We published an early version of the stock market paper in *Physica A* in 1992, and followed that with the version included here in 1997. The paper got considerable notice and went on to influence much further work on agent-based economics.

One other paper that was highly noticed came out in 1994, and this was my El Farol paper (included in this volume as Chapter 2). The idea had occurred to me at a bar in Santa Fe, El Farol. There was Irish music on Thursday nights and if the bar was not too full it was enjoyable, if the bar was crowded it was much less so. It occurred to me that if everyone predicted that many would come on a given night, they would not come, negating that forecast; and if everyone predicted few would come they would come, negating that forecast too. Rational forecasts—rational expectations—would be self-negating. There was no way to form properly functioning rational expectations. I was curious about what artificial agents might make of this situation and in 1993 I programmed it up and wrote a paper on it. The paper appeared in the *American Economic Review's Papers and Proceedings*, and economists didn't know at first what to make of it. But it caught the eye of Per Bak, the physicist who had originated the idea of self-organized criticality. He started to fax it to colleagues, and suddenly El Farol was well known in physics. Three years later, a game-theoretic version of the problem was introduced by the physicists Damien Challet and Yi-Cheng Zhang of the University of Freiburg as the Minority Game.[3] Now, several hundred papers later, both the Minority Game and El Farol have been heavily studied.

§

In 1997 my ideas took off in a different direction, one that wasn't directly related to Santa Fe's economics program. I became deeply interested in technology. The interest at first puzzled me. My early background was engineering, but still, this fascination with technology seemed to have nothing to do with my main interests in either economics or complexity. The interest had in fact been kindled years before, when I was exploring the idea of technologies competing for adoption. I had noticed that technologies—all the technologies I was looking at—had not come into being out of inspiration alone. They were all combinations of technologies that already existed. The laser printer had been put together from—was a combination of—a computer processor, a laser, and xerography: the processor would direct the laser to "paint" letters or images on a copier drum, and the rest was copying.

I had realized something else as well. In 1992 I had been exploring jet engines out of curiosity and I wondered why they had started off so simple yet within two or three decades had become so complicated. I had been learning C programming at the time, and it occurred to me that C programs were structured in basically the same way as jet engines, and as all technologies for that matter. They had a central functioning module, and other sub-modules hung off this to set it up properly and to manage it properly. Over time with a given technology, the central module could be squeezed to deliver more performance if sub-technologies were added to get past physical limits or to work

---

[3] D. Challet and Y-C. Zhang, "Emergence of Cooperation and Organization in an Evolutionary Game," *Physica A* 246: 407–418, 1997.

around problems, and so a technology would start off simple, but would add pieces and parts as it evolved. I wrote an essay in *Scientific American* in 1993 about why systems tended to elaborate.[4]

Somehow in all this I felt there was something general to say about technology—a general theory of technology was possible. I had started to read widely on technology, and decided I would study and know very well several particular technologies, somewhere between a dozen and twenty. In the end these included not just jet engines, but early radio, radar, steam engines, packet switching, the transistor, masers, computation, and even oddball "technologies" such as penicillin. Much of this study I did in St. John's College library in Santa Fe, some also in Xerox Parc where I was now working. I began to see common patterns emerging in how technologies had formed and come into being. They all captured and used phenomena: ultimately technologies are phenomena used for human purposes. And phenomena came along in families—the chemical ones, the electronic ones, the genomic ones—so that technologies formed into groups: industrial chemistry, electronics, biotechnology.

What became clear overall was that it wasn't just that individual technologies such as the jet engine evolved over their lifetimes. Technology—the whole collection of individual technologies—evolved in the sense that *all* technologies at any time, like all species, could trace a line of ancestry back to earlier technologies. But the base mechanism was not Darwinian. Novel technologies did not come into existence by the cumulation of small changes in earlier technologies: the jet engine certainly did not emerge from small changes in air piston engines. Novel technologies sprung from combining or integrating earlier technologies, albeit with human imagination and ingenuity. The result was a mechanism for evolution different from Darwin's. I called it Evolution by Combination, or Combinatorial Evolution.

This mechanism exists of course also in biological evolution. The major transitions in evolution are mostly combinations. Unicellular organisms became multicellular organisms by combination, and prokaryotes became eukaryotes by combination. But the occurrence of such events is rare, every few hundred million years at best. The day-to-day evolutionary mechanism in biology is Darwinian accumulation of small changes and differential selection of these. By contrast, in technology the standard evolutionary mechanism is combination, with Darwinian small changes following once a new technology exists.

I felt I now understood how technologies came into existence, and how the collection of technology evolved. I wanted to see if I could make such evolution work in the lab or on a computer. Around 2005 I was working at FXPAL, Fuji Xerox's think tank in Palo Alto, and I had met the computer scientist Wolfgang Polak. Could we create a computer experiment in which a soup of primitive technologies could be combined at random and the resulting combination—a potential new technology—tossed out if not useful but retained if useful and added to the soup for further combination? Would such a system creating successive integrations in this way bootstrap its way from simplicity to sophistication? We experimented with several systems, to no avail. Then we came across a beautiful paper by Richard Lenski in

[4] W. B. Arthur, "Why Do Things Become More Complex?" *Scientific American*, May 1993.

Nature,[5] where he and his colleagues had used the genetic algorithm to evolve digital circuits. Digital technologies seemed a natural medium to work in: if you combined two digital circuits you got another digital circuit; and the new circuit might do something useful or it might not.

Getting our experiment to work wasn't easy, but after a couple of months Polak got the system running and it began to "create" novel circuits from simple ones. Beginning with a soup of simple 2-bit *nand* circuits, the basic building block in digital circuits, we could press the return button to start the experiment and examine what had been created 20 hours later. We found circuits of all kinds. Elementary ones had formed first, then ones of intermediate complication such as a 4-bit *equals*, or 3-bit *less than*. By the end an 8-bit *exclusive-or*, 8-bit *and*, and an 8-bit *adder* had formed. Casually this may not seem that significant. But an 8-bit *adder* that works correctly (adding 8 bits of x to 8 bits of y to yield 9 bits for the result, z) is one of over $10^{177,554}$ circuits with 16 inputs and 9 outputs, and the chance of finding that randomly in 250,000 steps is negligible. Our successive integration process, of combining primitive building blocks to yield useful simple building blocks, and combining these again to create further building blocks, we realized was powerful. And actual technology had evolved in this way. It had bootstrapped its way from few technologies to many, and from primitive ones to highly complicated ones.

We published our experiment in *Complexity* but strange to say it was little noticed or commented on. My guess is that it fell between cracks. It wasn't biological evolution, it wasn't the genetic algorithm, it wasn't pure technology, and it wasn't economics. And the experiment didn't solve a particular problem. It yielded a toolbox or library of useful circuits, much like the library of useful functions that programming language designers provide. But it yielded this purely by evolution, and I found this a wonder. I have a degree in electrical engineering and Polak has one in computer science, but if you asked either of us to design an 8-bit *adder* we'd have to bone up on digital electronics and do this from scratch. Yet we had designed an algorithm that could design such circuits automatically by evolution. I found the idea of this remarkable, and of the papers assembled here this is one I am greatly taken by. It demonstrated evolution in action, and evolution by a different mechanism—by combination, or successive integration.

Somehow, I thought, all this had to fit with how an economy evolves, indeed how an economy forms in this first place. As I worked on technology, I realized that while the economy creates technology, more important, technology (the collective of technologies we use to meet our human needs) creates the economy. So the economy is not just a container for its technologies, it is an *expression* of them. As these technologies changed, and as whole new bodies of technology entered, the economy changed. It changed in what it did and how it did it, and it changed in the arrangements and institutions that fitted to the new ways of doing things. The economy, in other words, changed in structure.

I wrote all of these findings up in a book, *The Nature of Technology: What It Is and How It Evolves*, that appeared in 2009. It was well received, particularly by professional

---

[5] Lenski, R., C. Ofria, R. Pennock, and C. Adami, "The Evolutionary Origin of Complex Features, *Nature*, 423, 139–443, 2003.

engineers, and has gone into several languages. Some of the papers collected here were way stations on the path to this book and one was directly part of it. This work on technology took me 12 years from inception to completion, and I found it fascinating. Of particular wonder were the mechanisms by which the collective of technology evolved, and the realization that technology is a thing with considerable logical structure. Technology, I believe, studied in itself, is every bit as complicated and structured as the economy, or the legal system. And it is an object of considerable beauty.

§

The various lines of research that have made up this intellectual journey seemed to me at the time disparate and unconnected. But if I look back on them now, and on the work of other colleagues at Santa Fe and elsewhere, I see that what was forming from all this slowly and gradually was an approach to economics. I'd summed up my earlier understanding in a 1999 article in *Science*, [6] and the editor insisted I give this different approach a name. I called it "complexity economics." Looking back now, the features of complexity economics are clear. The economy is not necessarily in equilibrium; in fact it is usually in nonequilibrium. Agents are not all knowing and perfectly rational; they must make sense of the situations they are in and explore strategies as they do this. The economy is not given, not a simple container of its technologies; it forms from them and changes in structure as this happens. In this way the economy is organic, one layer forms on top of the previous ones; it is ever changing, it shows perpetual novelty; and structures within it appear, persist for a while, and melt back into it again. All this is not just a more poetic, humanistic view of the economy. It can be rigorously defined, and precisely probed and analyzed.

I'm often asked how this new approach fits with standard economics. Isn't it simply a variation of standard economics? And won't it be absorbed seamlessly into—"bolted on" to (in economist Richard Bronk's phrase)—the neoclassical framework? My answer on both counts is no. This different approach is not just the use of computers to do agent-based modeling, nor of adding a deeper understanding of technology change to endogenous growth models. It is economics done differently, economics based on different concerns—particularly on how nonequilibrium works—an economics where the problems are different and the very idea of a solution is also different.

One way to see this is to recognize that standard neoclassical economics comes out of a particular way of looking at the world. Neoclassical economics inherited the Enlightenment view that behind the seeming disorder of the world lay Order and Reason and Perfection. And it inherited much from the physics of the late 1800s, in particular the idea that large numbers of interacting identical elements could be analyzed collectively via simple mathematical equations. By the mid-1900s this led in turn to a hope that the core of economic theory could be captured in simple mathematically expressed principles and thereby axiomatized. Some parts, such as macroeconomics or the theory of institutions, might have to be left out, but the core of the field could be ordered and tamed, and reduced to mathematics.

---

[6] W. B. Arthur, "Complexity and the Economy," *Science*, April 2, 1999, 284, 107–109.

That program was at best only partially successful. It certainly cleaned up much of the sloppy logic that had passed as theory before, and led to a fresh respect for the workings of markets and for the inherent advantages of the capitalist system. But it also, I believe, led to a stiffness in thinking, to a righteousness in what was permitted as economic theory and what was not, and to a closedness to other ideas. Shut out were the effects on the economy of politics, of power, of class, of society, of fundamental uncertainty, and of formation and creation and development. In the end it could be argued that the program—at least the extreme hyper-rational version of it—failed. If it needed Popperian testing, its ideas were falsified spectacularly in 2008 and the aftermath of the financial meltdown. Nobody could claim that the market had lost half of its worth in a short time because companies had suddenly lost half of their usefulness; the companies were much as before. And nobody could claim either that unemployment rates of 20% and upward in some of the European economies were due to the suddenly changed preferences of the labor force; people wanted jobs just as before. In 2009 the *Economist* magazine noted wryly that Wall Street was not the only victim of the financial crash, standard neoclassical economics had collapsed along with it.

On reflection, it shouldn't be surprising that this highly purified form of economic thinking ran into difficulties. One lesson Western thought has had to learn slowly in modern times is that if we try hard enough to reduce anything to pure logic—for example if we try to pin down a final meaning of such concepts as Truth, or Being, or Life,or if we try to reduce some field such as philosophy or mathematics (or economics for that matter) to a narrow set of axioms—such attempts founder. The world cannot be reduced to pure logic and caged within it. Sooner or later it slips out to reveal its true messiness, and all such projects fail.

Slowly replacing the pure order of neoclassical economics is a new respect for reality, shared by many researchers in economics. Behavioral economics is one such approach being pressed forward; the psychology of markets is another. So too are theories of development that rely increasingly on understanding institutions and the workings of technology. And so too is the approach offered here which now has very many practitioners besides our initial group at Santa Fe.

One of things that has surprised me, and pleased me enormously, was that many of the "modern" themes in this approach fit well with ideas in Schumpeter, and Smith, and Mill, and Marx, and Keynes, and with the ideas of the institutionalists and political economists that followed. They too saw the economy as emerging from its technologies, as changing structurally, as not necessarily being in equilibrium, and with its decision-makers facing fundamental uncertainty. These connections have not yet been formally made; they are more like threads of thought that link these new ideas with ones discussed in the past. But they do show economics rediscovering some of what it lost. We are beginning to have a theoretical picture of the economy in formation and in nonequilibrium.

The papers collected here were written from when I first went to Santa Fe in 1987 until the present day. There is some overlap among them; but that is inevitable. They build on the ideas and work of many other people in economics, complexity, and other fields, in particular my Santa Fe colleagues John Holland, Stuart Kauffman, David Lane, and Richard Palmer. And they build also on the work of people not closely connected with our original Santa Fe group: in particular, Peter Allen, Rob Axtell, Josh Epstein, Alan Kirman, and Lee Tesfatsion, who all have contributed to this new approach. The papers here owe a great deal to the neoclassical formulation—that's after all what I was trained in. Many are full-blown analyses, others are essays. They are arranged more or less by theme rather than by when they were written, but the research papers are mostly in the first half of the book and the essays in the second. They can be read in any order and I encourage readers to follow whatever sequence appeals to them—and certainly to reach out to the work of others in this area.

Taken together, a theme or framework for thinking emerges from the papers here. In the place of agents in well-defined problems with well-defined probabilistic outcomes using perfect deductive reasoning and thereby arriving at an equilibrium, we have agents who must make sense out of the situation they face, who need to explore choices using whatever reasoning is at hand, and who live with and must adjust to an outcome that their very adjustments may cause perpetually to change.

In 1996 the historian of economic thought David Colander put forward an allegory in which economists a century ago stood at the base of two mountains whose peaks were hidden in the clouds. They wanted to climb the higher peak and had to choose one of the two. They chose the mountain of well-definedness and mathematical order, only to see when they had worked their way up and finally got above the clouds that the other mountain, the one of process and organicism, was far higher.

Many economists have started to climb that other mountain in the last few years. I will be interested to see what we will find along the way.

W. Brian Arthur
Palo Alto, California
January 2014