The Sciences of Complexity and "Origins of Order"

Stuart A. Kauffman

University of Pennsylvania

A new science, the science of complexity, is birthing. This science boldly promises to transform the biological and social sciences in the forthcoming century. My own book, *Origins of Order: Self Organization and Selection in Evolution*, (Kauffman, 1992), is at most one strand in this transformation. I feel deeply honored that Marjorie Grene undertook organizing a session at the Philosophy of Science meeting discussing *Origins*, and equally glad that Dick Burian, Bob Richardson and Rob Page have undertaken their reading of the manuscript and careful thoughts. In this article I shall characterize the book, but more importantly, set it in the broader context of the emerging sciences of complexity. Although the book is not yet out of Oxford press's quiet womb, my own thinking has moved beyond that which I had formulated even a half year ago. Meanwhile, in the broader scientific community, the interest in "complexity" is exploding.

A summary of my own evolving hunch is this: In a deep sense, E. coli and IBM know their respective worlds in the same way. Indeed, E. coli and IBM have each participated in the coevolution of entities which interact with and know one another. The laws which govern the emergence of knower and known, which govern the boundedly rational, optimally complex biological and social actors which have co-formed, lie at the core of the science of complexity. This new body of thought implies that the poised coherence, precarious, subject to avalanches of change, of our biological and social world is inevitable. Such systems, poised on the edge of chaos, are the natural talismen of adaptive order.

The history of this emerging paradigm conveniently begins with the "cybernetic" revolution in molecular biology wrought by the stunning discoveries in 1961 and 1963, by later Nobelists Francoise Jacob and Jacques Monod that genes in the humble bacterium, E. coli, literally turn one another on and off (Jacob and Monod, 1961, 1963). This discovery laid the foundation for the still sought solution of the problem of cellular differentiation in embryology. The embryo begins as a fertilized egg, the single cell zygote. Over the course of embryonic development in a human, this cell divides about 50 times, yielding the thousand trillion cells which form the newborn. The central mystery of developmental biology is that these trillions of cells become radically different from one another, some forming blood cells, others liver cells, still other nerve, gut, or gonadal cells. Previous work had shown that all the cells of a

PSA 1990, Volume 2, pp. 299-322

human body contain the same genetic instructions. How, then, could cells possibly differ so radically?

Jacob and Monods' discovery hinted the answer. If genes can turn one another on and off, then cell types differ because different genes are expressed in each cell type. Red blood cells have hemoglobin, immune cells synthesize antibody molecules and so forth. Each cell might be thought of as a kind of cybernetic system with complex genetic-molecular circuits orchestrating the activities of some 100,000 or more genes and their products. Different cell types then, in some profound sense, calculate how they should behave.

1. The Edge of Chaos

My own role in the birth of the sciences of complexity begins in the same years, when as a medical student, I asked an unusual, perhaps near unthinkable question. Can the vast, magnificent order seen in development conceivable arise as a spontaneous self organized property of complex genetic systems? Why "unthinkable"? It is, after all, not the answers which scientists uncover, but the strange magic lying behind the questions they pose to their world, knower and known, which is the true impulse driving profound conceptual transformation. Answers will be found, contrived, wrested, once the question is divined. Why "unthinkable"? Since Darwin, we have viewed organisms, in Jacob's phrase, as bricolage, tinkered together contraptions. Evolution, says Monod, is "chance caught on the wing". Lovely dicta, these, capturing the core of the Darwinian world view in which organisms are perfected by natural selection acting on random variations. The tinkerer is an opportunist, its natural artifacts are ad hoc accumulations of this and that, molecular Rube Goldbergs satisfying some spectrum of design constraints.

In the world view of bricolage, selection is the sole, or if not sole, the preeminent source of order. Further, if organisms are ad hoc solutions to design problems, there can be no deep theory of order in biology, only the careful dissection of the ultimately accidental machine and its ultimately accidental evolutionary history.

The genomic system linking the activity of thousands of genes stands at the summit of four billion years of an evolutionary process in which the specific genes, their regulatory intertwining and the molecular logic have all stumbled forward by random mutation and natural selection. Must selection have struggled against vast odds to create order? Or did that order lie to hand for selection's further molding? If the latter, then what a reordering of our view of life is mandated!

Order, in fact, lies to hand. Our intuitions have been wrong for thousands of years. We must, in fact, revise our view of life. Complex molecular regulatory networks inherently behave in two broad regimes separated by a third phase transition regime: The two broad regimes are chaotic and ordered. The phase transition zone between these two comprises a narrow third complex regime poised on the boundary of chaos. (Kauffman 1969, 1989; Fogleman-Soulie 1985; Derrida and Pomeau 1986; Langton 1991; Kauffman 1991, 1992). Twenty five years after the initial discovery of these regimes, a summary statement is that the genetic systems controlling ontogeny in mouse, man, bracken, fern, fly, bird, all appear to lie in the ordered regime near the edge of chaos. Four billion years of evolution in the capacity to adapt offers a putative answer: Complex adaptive systems achieve, in a lawlike way, the edge of chaos.

Tracing the history of this discovery, the discovery that extremely complex systems can exhibit "order for free", that our intuitions have been deeply wrong, begins with the intuition that even randomly "wired" molecular regulatory "circuits" with random "logic" would exhibit orderly behavior if each gene or molecular variable were controlled by only a few others. Notebooks from that period mix wire-dot diagrams of organic molecules serving as drugs with wire dot models of genetic circuitry. The intuition proved correct. Idealizing a gene as "on" or "off", it was possible by computer simulations to show that large systems with thousands of idealized genes behaved in orderly ways if each gene is directly controlled by only two other genes. Such systems spontaneously lie in the ordered regime. Networks with many inputs per gene lie in the chaotic regime. Real genomic systems have few molecular inputs per gene, reflecting the *specificity* of molecular binding, and use a *biased* class of logical rules, reflecting molecular simplicity, to control the on/off behavior of those genes. Constraint to the vast ensemble of possible genomic systems characterized by these two "local constraints" also inevitably yields genomic systems in the ordered regime. The perplexing, enigmatic, magical order of ontogeny may largely reflect large scale consequences of polymer chemistry.

Order for free. But more: The spontaneously ordered features of such systems parallels a host of ordered features seen in the ontogeny of mouse, man, bracken, fern, fly, bird. A "cell type" becomes a stable recurrent pattern of gene expression, an "attractor" in the jargon of mathematics, where an attractor, like a whirlpool, is a region in the state space of all the possible patterns of gene activities to which the system flows and remains. In the spontaneously ordered regime, such cell type attractors are inherently small, stable, and few, implying that the cell types of an organism traverse their recurrent patterns of gene expression in hours not eons, that homeostasis, Claude Bernard's conceptual child, lies inevitably available for selection to mold, and, remarkably, that it should be possible to predict the number of cell types, each a whirlpool attractor in the genomic repertoire, in an organism. Bacteria harbor one to two cell types, yeast three, ferns and bracken some dozen, man about two hundred and fifty. Thus, as the number of genes, called genomic complexity, increases, the number of cell types increases. Plotting cell types against genomic complexity, one finds that the number of cell types increases as a square root function of the number of genes. And, outrageously, the number of whirlpool attractors in model genomic systems in the ordered regime also increase as a square root function of the number of genes. Man, with about 100,000 genes should have three hundred seventy cell types, close to two hundred and fifty. A simple alternative theory would predict billions of cell types.

Bacteria, yeast, ferns, and man, members of different phyla, have no common ancestor for the past 600 million years or more. Has selection struggled for 600 million years to achieve a square root relation between genomic complexity and number of cell types? Or is this order for free so deeply bound into the roots of biological organization that selection cannot *avoid this order*? But if the latter, then *selection is not the sole source of order in biology*. Then Darwinism must be extended to embrace self organization and selection.

The pattern of questions posed here is novel in biology since Darwin. In the NeoDarwinian world view, where organisms are ad hoc solutions to design problems, the answers lie in the specific details wrought by ceaseless selection. In contrast, the explanatory approach offered by the new analysis rests on examining the statistically typical, or generic, properties of an entire class, or "ensemble" of systems all sharing known local features of genomic systems. If the typical, generic, features of ensemble members corresponds to that seen in organisms, then explanation of those features emphatically *does not* rest in the details. It rests in the general laws governing the typical features of the ensemble as a whole. Thus an "ensemble" theory is a new kind of statistical mechanics. It predicts that the typical properties of members of the ensemble will be found in organisms. Where true, it bodes a physics of biology.

Not only a physics of biology, but beyond, such a new statistical mechanics demands a new pattern of thinking with respect to biological and even cultural evolution: Self organization, yes, aplenty. But selection, or its analogues such as profitability, is always acting. We have no theory in physics, chemistry, biology, or beyond which marries self organization and selection. The marriage consecrates a new view of life.

But two other failures of Darwin, genius that he was, must strike us. How do organisms, or other complex entities, manage to adapt and learn? That is, what are the conditions of "evolvability". Second, how do complex systems coordinate behavior, and more deeply, why are adaptive systems so often complex?

Consider "evolvability" first. Darwin supposed that organisms evolve by the *successive accumulation of useful random variations*. Try it with a standard computer program. Mutate the code, scramble the order of instructions, and try to "evolve" a program calculating some complex function. If you do not chuckle, you should. Computer programs of the familiar type are not readily "evolvable". Indeed the more compact the code, the more lacking in redundancy, the more sensitive it is to each minor variation. Optimally condensed codes are, perversely, minimally evolvable. Yet the genome is a kind of molecular computer, and clearly has succeeded in evolving. But this implies something very deep: Selection must achieve the kinds of systems which are able to adapt. That capacity is not Godgiven, it is a success.

If the capacity to evolve must itself evolve, then the new sciences of complexity seeking the laws governing complex adapting systems must discover the laws governing the emergence and character of systems which can themselves adapt by accumulation of successive useful variations.

But systems poised in the ordered regime near its boundary are precisely those which can, in fact, evolve by successive minor variations. The behavior of systems in the chaotic regime are so drastically altered by any minor variation in structure or logic that they cannot accumulate useful variations. Conversely, systems deep in the ordered regime are changed so slightly by minor variations that they adapt too slowly to an environment which may sometimes alter catastrophically. Evolution of the capacity to adapt would be expected, then, to achieve poised systems.

How can complex systems coordinate behavior? Again, complex adaptive entities achieve the edge of chaos because such systems can coordinate the most complex behavior there. Deep in the chaotic regime, alteration in the activity of any element in the system unleashes an avalanche of changes, or "damage", which propagates throughout most of the system (Stauffer 1987). Such spreading damage is equivalent to the "butterfly effect" or sensitivity to initial conditions typical of chaotic systems. The butterfly in Rio changes the weather in Chicago. Crosscurrents of such avalanches unleashed from different elements means that behavior is not controlable. Conversely, deep in the ordered regime, alteration at one point in the system only alters the behavior of a few neighboring elements. Signals cannot propagate widely throughout the system. Thus, control of complex behavior cannot be achieved. Just at the boundary between order and chaos, the most complex behavior can be achieved.

Finally, computer simulations suggest that natural selection or its analogues actually do achieve the edge of chaos. This third regime, poised between the broad ordered regime and the vast chaotic regime, is razorblade thin in the space of systems.

Absent other forces, randomly assembled systems will lie in the ordered or chaotic regimes. But let such systems play games with one another, winning and losing as each system carries out some behavior with respect to the others, and let the structure and logic of each system evolve by mutation and selection, and, lo, systems do actually adapt to the edge of chaos! No minor point this: Evolution itself brings complex systems, when they must adapt to the actions of other, to an internal structure and logic poised between order and chaos, (Kauffman 1991).

We are lead to a bold hypothesis: Complex adaptive systems achieve the edge of chaos.

The story of the "edge of chaos" is stronger, the implications more surprising. Organisms, economic entities, nations, do not evolve, they *coevolve*. Almost miraculously, coevolving systems, too, mutually achieve the poised edge of chaos. The sticky tongue of the frog alters the fitness of the fly, and deforms its fitness landscapes that is, what changes in what phenotypic directions improve its chance of survival. But so too in technological evolution. The automobile replaced the horse. With the automobile came paved roads, gas stations hence a petrolium industry and war in the Gulf, traffic lights, traffic courts, and motels. With the horse went stables, the smithy, and the pony express. New goods and services alter the economic landscape. Coevolution is a story of coupled deforming "fitness landscapes". The outcome depends jointly on how much my landscape is deformed when you make an adaptive move, and how rapidly I can respond by changing "phenotype".

Are there laws governing coevolution? And how might they relate to the edge of chaos? In startling ways. Coevolution, due to a selective "metadynamics" tuning the structure of fitness landscapes and couplings between them, may typically reach the edge of chaos (Kauffman 1992). E.coli and IBM not only "play" games with the other entities with which they coevolve. Each also participates in the very *definition or form* of the game. It is we who create the world we mutually inhabit and in which we struggle to survive. In models where players can "tune" the mutual game even as they play, or coevolve, according to the game existing at any period, the entire system moves to the edge of chaos. This surprising result, if general, is of paramount importance. A simple view of it is the following: Entities control a kind of "membrane" or boundary separating inside from outside. In a kind of surface to volume way, if the surface of each system is small compared to its volume it is rather insensitive to alterations in the behaviors of other entities. That is, adaptive moves by other partners do not drastically deform one partner's fitness landscape. Conversely, the ruggedness of the adaptive landscape of each player as it changes its "genotype" depends upon how dramatically its behavior deforms as its genotype alters. In turn this depends upon whether the adapting system is itself in the ordered, chaotic, or boundary regime. If in the ordered, the system itself adapts on a smooth landscape. In the chaotic regime the system adapts on a very rugged landscape. In the boundary regime the system adapts on a landscape of intermediate ruggedness, smooth in some directions of "genotype" change, rugged in other directions. Thus, both the ruggedness of one's own fitness landscape and how badly that landscape is deformed by moves of one's coevolving partners are *themselves possible objects of a selective "metadynamics"*. Under this selective metadynamics, tuning landscape structure and susceptibility, model coevolving systems which mutually know and interact with one another actually reach the edge of chaos. Here, under most circumstances, most entity optimizes fitness, or payoff, by remaining the same. Most of the ecosystem is frozen into a percolating Nash equilibrium, while coevolutionary changes propagate in local unfrozen islands within the ecosystem. More generally, alterations in circumstances send avalanches of changed optimal strategies propagating through the coevolving system. At the edge of chaos the size distributions of those avalanches approach a power law, with many small avalanches and few large ones. During such coevolutionary avalanches, affected players would be expected to fall transiently to low fitness, hence might go extinct. Remarkably, this size distribution comes close to fitting the size distribution of extinction events in the record. At a minimum, a distribution of avalanche sizes from a common size small cause tells us that small and large extinction events may reflect endogenous features of coevolving systems more than the size of the meteor which struck.

The implications are mini-Gaia. As if by an invisible hand, coevolving complex entitites may mutually attain the poised boundary between order and chaos. Here, mean sustained payoff, or fitness, or profit, is optimized. But here avalanches of change on all length scales can propagate through the poiised system. Neither Sisyphis, forever pushing the punishing load, nor fixed unchanging and frozen, instead E.coli and its neighbors, IBM and its neighbors, even nation states in their collective dance of power, may attain a precarious poised complex adaptive state. The evolution of complex adaptive entities itself appears lawful. How far we come from Darwin's genius.

This strand in the birth of complexity theory, here spun, has its history. The first stages were set in the mid 1960s by the discovery of spontaneous order, as well as the expected chaos, in complex genomic systems. The discovery was not without attention among scientists the day. Warren McCulloch, patriarch of cybernetics, author with Pitts of "The Logical Calculus of Ideas Imminent in the Mind", step-child of Bertrand Russell's logical atomism, and ancestor to todays neural connectionist tony flowering, invited me to share his home with his remarkable wife Rook. "In pine tar is. In oak none is. In mud eels are. In clay none are", sang this poet of neural circuitry, demonstrating by dint of a minor Scots accent that no hearer could unscramble four simple declarative sentences. Mind, complex, could fail to classify. "All Cambridge excited about your work", wrote McCulloch to this medical student who, thrilled, was yet to decode Warren's style.

Yet the time was not ripe. McCulloch had said twenty years would elapse before biologists took serious note. He was right, almost to the hour. And for good reason had he made his prediction. The late 1960s witnessed the blunderbuss wonderful explosion of molecular biology. Enough, far more than enough, to thrill to the discovery of the real molecular details: How a gene is transcribed to RNA, translated to protein, acts on its neighbors. What is the local logic of a bacterial genetic circuit controlling metabolism of lactose? Of a bacterial virus, or phage? What of the genes in a higher organism like the heralded but diminutive fruit fly? What of mouse and man? Enveloped by the Darwinian world view, whose truths run deep, held in tight thrall by the certainty that the order in organisms resides in the well wrought details of construction and design, details inevitably ad hoc by virtue of their tinkered origins in the wasteland of chance, molecular biologists had no use for heady, arcane, abstract ensemble theories. The birth of complexity theory, or this strand of it, though noted, received no sustaining passion from its intended audience.

Twenty years, indeed. Rebirth of this strand was midwifed by the physicists. An analogue ensemble theory, called "spin glasses", had been developed starting in the mid 1970s by solid state physicists such as Philip Anderson, Scott Kirkpatrick, Bernard Derrida, Gerard Toulouse, were struggling with an odd kind of dilute magnet material. Unlike the familiar ferromagnet, captured in the famous Ising model, where magnetic spins like to orient in the same direction as their neighboring spins, hence the magnetized state with all spins oriented in the same direction arises, in these bewildering spin glasses, adjacent spins might like to orient in the same or in the opposite direction, depending sinusoidally on the distance between the spins. What a mess. Edwards and Anderson started an industry among their brethren, and legitimized the new class of ensemble theories, by building mathematical models of spin glasses on two or three dimensional lattices,. Here each vertex houses a spin. But, to capture the bizarre logic of their magnetic materials, Edwards and Anderson assumed that each adjacent pair of spins "chose", once and forever, whether they wanted to point in the same or opposite direction, and how much they cared, given by an energy for that bond. Such messy models meant two major things. First, since couplings are assigned at random, any one model spin glass is a member of a vast ensemble governed by the same statistics. This is an ensemble theory averaging, not over the states of one system as in the familiar statistical mechanics of gases, but over billions of systems in the same ensemble. One seeks and characterizes the typical, or generic features of these systems. Second, such systems have tortuous and rugged "energy landscapes". This is due to "frustration". Consider four spins around a square, where three pairs wish to point in the same direction, the fourth does not. All cannot be satisfied. Each configuration of the many spins in the lattice of a spin glass has a total energy. The distribution of energies over the configurations is the energy landscape, the analogue of a fitness landscape. Frustration implies that the landscape is rugged and multipeaked.

Later, the structures of these spin glass landscapes would provide new models of molecular evolution over rugged multipeaked fitness landscapes. Molecular evolution turns out to be much like an electron bouncing on a complex potential surface at a small temperature. At too low a temperature, the electron remains trapped in poor potential wells. At too high a temperature, the electron bounces all over the potential surface and has a high, unhappy, average energy. On any defined time scale, energy is minimized at a specific fixed temperature at which the electron is just "melting" out over the energy landscape, sliding gracefully over low saddles in the surface separating wells such that it finds good potential wells rather easily, then does not hop out of them too rapidly. The analogue in molecular evolution or other biological evolution over a fixed fitness landscape, or one deforming at a given mean rate, is to tune the parameters of adaptive search over the space such that an adapting population is just "melting" out of local regions of the space. Again: The edge of Chaos!

By 1985 many of the physicists had tired of their spin glasses. Some turned to models of neural networks, sired by McCulloch, where neurons turn one another on and off rather like genes, or like spins for that matter. Hopfield found further fame by modeling parallel processing neural networks as spin systems, (Hopfield 1982). Attractors of such networks, rather than modeling cell types as I had suggested, were taken to model memories. Each memory was an attractor. Memories were content addressable, meaning that if the network were begun in the "basin of attraction" drained by one whirlpool attractor, the system would flow to that attractor. Partial data, corresponding to an initial state in a basin of attraction but not on the attractor itself, could be reconstructed to the stored memory. (All scientists regret the article not written. Jack Cowan and I had sketched an article in 1970 arguing against the logical atomism implicit in McCulloch and Pitts, an atomism melted by Wittgenstein's Investigations. In contrast, we wanted to suggest that concepts were attractors in neural networks, hence a collective integrated activity. From Wittgenstein we knew that language games are not reducible to one another, law to human action to physical phenomena. We wanted to argue that new concepts, new language games, arose by bifurcations yielding new attractors in the integrated activity of coupled neurons. One such attractor would not be reducible in any obvious way to another attractor. Grandmother cells be damned, concepts are collective properties.) Toulouse, brilliant as Hopfield, followed with other spin glass like models whose basins of attraction were, he said, more like French than English gardens. Many have followed, to the field's flowering.

Not all the physicists who tired of spin glasses turned to neurobiology. In the way of these things, French physicist Gerard Weishbuch was romantically involved with French mathematician Francoise Fogleman-Soulie. Francoise chose, as her thesis topic, the still poorly understood order found in "Kauffman nets", (Fogleman-Soulie 1985). Many theorems followed. Gerard's interest extended from Francoise and spin glasses to this strange hint of order for free. Summers in Jerusalem and Haddasah hospital with Henri Atlan, doctor, theoretical biologist, author of Crystal and Smoke with its search for order and adaptability, led to more results. Put these bizarre genetic networks on lattices, where any good problem resides. See the order. Scale parameters. Find phase transitions and the scaling laws of critical exponents. A new world to a biologist. And Gerard shared an office with Bernard Derrida, nephew of deconstructionist Jacques. Bernard looked at at these "Kauffman nets", the name is due to Derrida, and leaped to an insight no biologist would ever dare. Let the network be randomly rewired at each moment, creating an "annealed" model. Theorem followed theorem. No genome dances so madhatterly. But the mathematics can. Phase transition assured. Order for free in networks of low connectivity. Analysis of sizes of basins of attraction, and of overlaps between attractors, (Derrida and Pomeau 1986). I lost a bottle of wine to Derrida, shared over dinner, on the first theorem.

Even I chimed in with a few theorems here and there: a mean field approach to attractors, the existence of a connected set of elements which are "frozen" and do not twinkle on and off, that spans or percolates across the system. This frozen component, leaving behind isolated twinkling islands, is the hallmark of order. The phase transition to chaos occurs, as parameters alter, when the frozen component "melts", and the twinkling islands merge into an unfrozen, twinkling, percolating sea, leaving behind small isolated frozen islands. The third, complex regime, the boundary between order and chaos, arises when the twinkling connected, percolating sea is just breaking up into isolated islands. Avalanches of changes due to perturbations, which only propagate in the twinkling unfrozen sea, show a characteristic "power law" distribution at the phase transition, with many small avalanches and a few enormous ones, (Kauffman 1989).

Now the reader can see why systems on the boundary between order and chaos can carry out the most complex tasks, adapt in the most facile fashion. Now too, I hope, you can see the intrigue.at the possibility that complex adaptive systems achieve the edge of chaos in their internal structure, but may also coevolve in a selective metadynamics to achieve the edge of chaos in the ecosystem of the mutual games they play! The edge of chaos may be a major organizing principle governing the evolution and coevolution of complex adaptive systems.

Other themes, again spawned by physicists, arose in America, and lead quasi-independently, quasi-conversing, to the growth of interest in complexity. "Kauffman nets", where the wiring diagram among "genes" or binary elements, is random, and the logic governing each element is randomly assigned, hence differs for different "genes", are versions of a mathematical structure called "cellular automata". Cellular automata were invented by von Neuman, whose overwhelming early work, here and on the existence of self reproducing automata, filters down through much that fol lows. The simplest cellular automata are lines or rings of on/off sites, each governed by the same logical rule which specifies its next activity, on or off, as a function of its own current state and those of its neighbors to a radius, r. Enter young Stephen Wolfram, quick, mercurial, entreprenurial. The youngest MacArthur Fellow, Wolfram had begun publishing in high energy physics at age 16. While a graduate student at Cal Tech, he earned the mixed admiration and emnity of his elders by inventing computer code to carry out complex mathematical calculations. Cal Tech did not mind his mind. It minded his marketing the products of his mind. Never mind. Thesis done, Wolfram packed off to the Institute for Advanced Study and fell to the analysis of cellular automata. He amazed his audiences. The world of oddball mathematicians, computer scientists, wayward physicists, biologists soon twiddled with CA rules. Four classes of behavior emerged, stable, periodic, and chaotic, of course. And between them, on the edge between order and chaos, capable of complex computation, perhaps universal computation? A third "complex class". Among the most famous of these CA rules is Conway's "Game of Life", provable capable of universal computation, demonstrably capable of capturing gigabits of memory and gigaseconds of time among amateurs and professionals world wide. The game of life, like true life itself according to our bold hypothesis, also lies at the edge of chaos.

Paralleling Derrida is the lineage flowing from Chris Langton. Langton, a computer scientist and physicist, elder graduate student, survivor of early hang gliding and an accident relieving him of most unbroken bone structure in his mid-twenties body, thought he could improve on von Neuman. He invented a simple self reproducing automaton and littered computer screens from Los Alamos to wherever. Then Langton, following von Neuman again, and fired up by Wolfram, began playing with cellular automata. Where I had shown that the transition from order to chaos was tuned by tuning the number of inputs per "gene" from 2 to many, Langton reinvented Derrida's approach. Derrida, like Langton after him, in turn reinvented a classification of logical rules first promulgated by Crayton Walker. This classification marks the bias, P, towards the active, or inactive state, over all combinations of activities of the inputs to an element. Derrida had shown that the phase transition occurred at a critical value of this bias, Pc. At that bias, frozen components emerge. Langton found the same phase transition, but measured in a different way to focus on how complex a computation might be carried out in such a network. This complexity, measured as mutual information, or what one can predict about the next activity of one site given the activity of another site, is maximized at the phase transition (Langton 1991).

The poised edge reappears, like a new second law of thermodynamics, everywhere hinted, but, without Carnot, not yet clearly articulated, in the recent work of physicist Jim Crutchfield. "Symbolic dynamics" is a clever new tool used to think about complex dynamical systems. Imagine a simple system such as a pendulum. As its swings back and forth, it crosses the midpoint where it hangs straight down. Use a 1 to denote times when the pendulum is to the left of the midpoint, and 0 to denote times when the pendulum swings to the right. Evidently, the periodic pendulum gives rise to an alternating sequence of 1 and 0 values. Such a symbol sequence records the dynamics of the pendulum by breaking its state space into a finite number of regions, here two, and labeling each region with a symbol. The flow of the system gives rise to a symbol sequence. Theorems demonstrate that, with optimally chosen boundaries between the regions, here the midpoint, the main features of the dynamics of the real pendulum can be reconstructed from the symbol sequence. For a periodic process, the symbol sequence is dull. But link several pendulums together with weak springs and again denote the behavior of one pendulum by 1 and 0 symbols. Now the motion of each pendulum is influenced by all the others in very complex ways. The symbol sequence is correspondingly complex. The next step is to realize that any symbol sequence can be generated as the output of a finite automaton, a more or less complex "neural" or "genetic" network of on off elements. Further, theorems assure us that for any such symbol sequence, the smallest, or minimal automaton, with the minimal number of elements and internal states, can be found. Thus, the number of elements, or states, of such a system is a measure of the complexity of the symbol sequence. And now the wonderful surprise. The same three phases, ordered, chaotic, and complex, are found again. That is, such automata, like Kauffman nets and neural nets, harbor the same generic behaviors. And, as you will now suspect, the complex regime again corresponds to the most complex symbol sequences, which in turn arise in dynamical systems themselves on the boundary between order and chaos.

If one had to formulate, still poorly articulated, the general law of adaptation in complex systems, it might be this: Life adapts to the edge of chaos.

2. The Origin of Life and its Progeny

This story, the story of the boundary between order an chaos achieved by complex coevolving systems, is but half the emerging tale. The second voice tells of the origin of life itself, a story both testable and, I hope, true, a story implying vast stores of novel drugs, vaccines, universal enzymatic tool boxes, a story latent with the telling of technological and cultural evolution, of bounded rationality, the coemergence of knower and known, hence at last, of telling whether E. coli and IBM do, in fact, know their worlds, the worlds they themselves created, in the same deep way.

Life is held a miracle, God's breath on the still world, yet cannot be. Too much the miracle, then we were not here. There must be a viewpoint, a place to stand, from which the emergence of life is explicable, not as a rare untoward happening, but as expected, perhaps inevitable. In the common view, life originated as a self reproducing polymer such as RNA, whose self complementary structure, since Watson and Crick remarked with uncertain modesty, suggests its mode of reproduction, has loomed the obvious candidate urbeast to all but the stubborn. Yet stubbornly resistant to test, to birthing in vitro is this supposed simplest molecule of life. No worker has yet succeeded in getting one single stranded RNA to line up the complementary free nucleotides, link them together to form the second strand, melt them apart, then repeat the cycle. The closest approach shows that a polyC polyG strand, richer in C than G, can in fact line up its complementary strand. Malevolently, the newly formed template is richer in G than C, and fails, utterly, to act as a facile template on its own. Alas.

Workers attached to the logic of molecular complementarity are now focusing effort on polymers other than RNA, polymers plausibly formed in the prebiotic environment, which might dance the still sought dance. Others, properly entranced with the fact that RNA can act as an enzyme, called a ribozyme, cleaving and ligating RNA sequences apart and together, seek a ribozyme which can glide along a second RNA, serving as a template that has lined up its nucleotide complements, and zipper them together. Such an ribozyme would be a ribozyme polymerase, able to copy any RNA molecule, including itself. Beautiful indeed. And perhaps such a molecule occurred at curtain-rise or early in the first Act. But consider this: A free living organism, even the simplest bacterium, links the synthesis and degradation of some thousands of molecules in the complex molecular traffic of metabolism to the reproduction of the cell itself. Were one to begin with the RNA urbeast, a nude gene, how might it evolve? How might it gather about itself the clothing of metabolism?

There is an alternative approach which states that life arises as a nearly inevitable phase transition in complex chemical systems. Life formed by the emergence of a collectively autocatalytic system of polymers and simple chemical species.

Picture, strangely, ten thousand buttons scattered on the floor. Begin to connect these at random with red threads. Every now and then, hoist a button and count how many buttons you can lift with it off the floor. Such a connected collection is call a "component" in a "random graph". A random graph is just a bunch of buttons connected at random by a bunch of threads. More formally, it is a set of N nodes connected at random by E edge. Random graphs undergo surprising phase transitions. Consider the ratio of E/N, or threads divided by buttons. When E/N is small, say .1, any button is connected directly or indirectly to only a few other buttons. But when E/N passes 0.5, so there are half as many threads as buttons, a phase transition has occurred. If a button is picked up, very many other buttons are picked up with it. In short, a "giant component" has formed in the random graph in which most buttons are directly or indirectly connected with one another. In short, connect enough nodes and a connected web "crystalizes".

Now life. Proteins and RNA molecules are linear polymers build by assembling a subset of monomers, twenty types in proteins, four in RNA. Consider the set of polymers up to some length, M, say 10. As M increases the number of types of polymers increases exponentially, for example there are 20 M proteins of length M. This is a familiar thought. The rest are not. The simplest reaction among two polymers consists in gluing them together. Such reactions are reversible, so the converse reaction is simply cleaving a polymer into two shorter polymers. Now count the number of such reactions among the many polymers up to length M. A simple consequence of the combinatorial character of polymers is that there are many more reactions linking the polymers than there are polymers. For example, a polymer length M can be formed in M 1 ways by gluing shorter fragments comprising that polymer. Indeed, as M increases, the ratio of reactions among the polymers to polymers is about M, hence increases as M increases. Picture such reactions as black, not red, threads running from the two smaller fragments to a small square box, then to the larger polymer made of them. Any such triad of black threads denotes a possible reaction among the polymers; the box, assigned a unique number, labels the reaction itself. The collection of all such triads is the chemical reaction graph among them. As the length of the longest polymer under consideration, M, increases, the web of black triads among these grows richer and richer. The system is rich with crosslinked reactions.

Life is an autocatalytic process where the system synthesizes itself from simple building blocks. Thus, in order to investigate the conditions under which such a autocatalytic system might spontaneously form, assume that no reaction actually occurs unless that reaction is catalyzed by some molecule. The next step notes that protein and RNA polymers can in fact catalyse reactions cleaving and ligating proteins and RNA polymers: trypsin in your gut after dinner digesting steak, or ribozyme ligating RNA sequences. Build a theory showing the probability that any given polymer catalyzes any given reaction. A simple hypothesis is that each polymer has a fixed chance, say one in a billion, to catalyze each reaction. No such theory can now be accurate, but this hardly matters. The conclusion is at hand, and insensitive to the details. Ask each polymer in the system, according to your theory, whether it catalyzes each possible reaction. If "yes", color the corresponding reaction triad "red", and note down which polymer catalyzed that reaction. Ask this question of all polymers for each reaction. Then some fraction of the black triads have become red. The red triads are the catalyzed reactions in the chemical reaction graph. But such a catalyzed reaction graph undergoes the button thread phase transition. When enough reactions are catalyzed, a vast web of polymers are linked by catalyzed reactions. Since the ratio of reactions to polymers increases with M, at some point as M increases at least one reaction per polymer is catalyzed by some polymer. The giant component crystalizes. An autocatalytic set which collectively catalyzes its own formation lies hovering in the now pregnant chemical soup. A self reproducing chemical system, daughter of chance and number, swarms into existence, a connected collectively autocatalytic metabolism. No nude gene, life emerged whole at the outset.

I found this theory in 1971. Even less than order for free in model genomic systems did this theory find favor. Stuart Rice, colleague, senior chemist, member of the National Academy of Science asked, "What for?" Alas again. When famous older scientists say something warrants the effort, rejoice. When famous older scientists are dismissive, beware. I turned to developmental genetics and pattern formation, the beauty of Alan Turing's theory of pattern formation by the establishment of chemical waves, the quixotic character of homeotic mutants in the fruit fly, Drosophila melanogaster, where eyes convert to wings, antennas to legs, and heads to genitalia. Fascinating disorders, these, called metaplasias, whose battered sparse logic hinted the logic of developmental circuits. But experimental developmental genetics, even twelve years and surgery on ten thousand embryos, is not the central thread of the story.

In 1983 interest in serious theories of the origin of life was rekindled. In 1971 and the ensuing decade, Nobelist Manfred Eigen, together with theoretical chemist Peter Schuster, developed a well formulated, careful model of the origin of life, called the "hypercycle". In this theory, the authors begin by assuming that short nude RNA sequences can replicate themselves. The hooker is this: During such replication, errors are made. The wrong nucleotide may be incorporated at any site. Eigen and Schuster showed that an error catastrophe occurs when RNA sequences become too long for any fixed error rate. The RNA population "melts" over RNA sequence space, hence all information accumulated within the "best" RNA sequence, culled by natural selection, is lost. The "hypercycle" is a clever answer to this devastation: Assume a set of different short RNA molecules, each able to replicate itself. Now assume that these different RNA molecules are arranged in a control cycle, such that RNA 1 helps RNA 2 to replicate, RNA 2 helps RNA 3, and so on until RNA N closes the loop by helping RNA 1. Such a loop is a hypercycle, "hyper" because each RNA itself is a tiny cycle of two complementary strands which copy one another. The hypercycle is, in fact, a coevolving molecular society. Each RNA species coevolves in company with its peers. This model has been studied in detail, and has strengths and weakness. Not the least of the latter is that no RNA sequence can yet replicate itself.

But other voices were lifted, from the most intelligent minds. Freeman Dyson, of the Institute of Advanced Studies, an elegant scientist and author of lyric books such as "Disturbing the Universe", suggested, in Origins of Life, that life arose as a phase transition in complex systems of proteins. Philip Anderson, with Daniel Stein, and Rothkar, borrowed from spin-glass theory to suggest that a collection of template replicating RNA molecules with overlapping ends and complex fitness functions governing their survival might give rise to many possible self reproducing sequences.

Lives in science have their peculiar romance. I heard of these approaches at a conference in India. Central India, Madya Pradesh, sweats with the sweet smell of the poor cooking over fires of dried buffalo dung. The spiritual character of India allows one to speak of the origin of life with colleagues such as Humberto Maturana, riding in disrepair except for his glasses and clear thoughts, in a bus of even greater disrepair among the buffalo herds to Sanchi, early Buddest shrine. The Budda at the west portal, thirteen hundred years old, ineffably young, invited only a gentle kiss from the foreigners in time, space, culture. Dyson's and Anderson's approaches appeared flawed. Dyson had assumed his conclusion, hidden in assumption 7. Life as an autocatalytic crystalization was trivially present in his model, slipped in by hand, not accounted for as a deeply emergent property of chemistry. And Anderson, overwhelmingly insightful, proposed nothing truly deep not already resting on RNA self complementarity. The romance continues with a flurry of theorems and lemmas, simple to a real mathematician.

This hiccup of creativity, I hoped, warranted investigation. Doyne Farmer, young physicist at Los Alamos, and his childhood friend Norman Packard, and I began collaborating to build detailed computer simulations of such autocatalytic polymer systems. Six years later, a Ph.D. thesis by Richard Bagley later, it is clear that the initial intuitions were fundamentally correct: In principle complex systems of polymers can become collectively self reproducing. The routes to life are not twisted backalleys of thermodynamic improbability, but broad boulevards of combinatorial inevitability.

If this new view of the crystalization of life as a phase transition is correct, then it should soon be possible to create actual self reproducing polymer systems, presumably of RNA or proteins, in the laboratory. Experiments, even now, utilizing very complex libraries of RNA molecules to search for autocatalytic sets are underway in a few laboratories.

If not since Darwin, then since Weisman's doctrine of the germ plasm was reduced to molecular detail by discovery of the genetic role of chromosomes, biologist have believed that evolution via mutation and selection virtually requires a stable genetic material as the store of heritable information. But mathematical analysis of autocatalytic polymer systems belies this conviction. Such systems can evolve to form new systems. Thus, contrary to Richard Dawkin's thesis in "The Selfish Gene", biological evolution does not, in principle, demand self-replicating genes at the base (Dawkins 1976). Life can emerge and evolve without a genome. Heresy, perhaps? Perhaps.

Many and unexpected are the children of invention. Autocatalytic polymer sets have begotten an entire new approach to complexity.

The starting point is obvious. An autocatalytic polymer set is a functional integrated whole. Given such a set, it is clear that one can naturally define the function of any given polymer in the set with respect to the capacity of the set to reproduce itself. Lethal mutants exist, for if a given polymer is removed, or a given foodstuff deleted, the set may fail to reproduce itself. Ecological interactions among coevolving autocatalytic sets lie to hand. A polymer from one such set injected into a second such set may block a specific reaction step and "kill" the second autocatalytic set. Coevolution of such sets, perhaps bounded by membranes, must inevitably reveal how such systems "know" one another, build internal models of one another, and cope with one another. Models of the evolution of knower and known lay over the conceptual horizon.

Walter Fontana, graduate student of Peter Schuster, came to the Santa Fe Institute, and Los Alamos. Fontana had worked with John McCaskill, himself an able young physicist collaborating with Eigen at the Max Planck Institute in Gottingen. McCaskill dreamt of polymers, not as chemicals, but as Turing machine computer programs and tapes. One polymer, the computer, would act on another polymer, the tape, and "compute" the result, yielding a new polymer. Fontana was entranced. But he also found the autocatalytic story appealing. Necessarily, he invented "Algorithmic Chemistry". Necessarily, he named his creation "Alchemy", (Fontana 1991).

Alchemy is based on a language for universal computation called the lambda calculus. Here almost any binary symbol string is a legitimate "program" which can act on almost any binary symbol string as an input to compute an output binary symbol string. Fontana created a "Turing gas" in which an initial stock of symbol strings randomly encounter one another in a "chemostat" and may or may not interact to yield symbol strings. To maintain the analogue of selection, Fontana requires that a fixed total number of symbol string polymers be maintained in the chemostat. At each moment, if the number of symbol strings grows above the maximum allowed, some randomly strings are lost from the system.

Autocatalytic sets emerge again! Fontana finds two types. In one, a symbol string which copies itself emerges. This "polymerase" takes over the whole system. In the second, collectively autocatalytic sets emerge in which each symbol string is made by some other string or strings, but none copies itself. Such systems can then evolve in symbol string space evolution without a genome.

Fontana had broken the bottleneck. Another formulation of much the same ideas, which I am now using, sees interactions among symbol strings creating symbol strings and carrying out a "grammar". Work on disordered networks, work which exhibited the three broad phases, ordered, chaotic and complex, drove forward based on the intuition that order and comprehensibility would emerge by finding the generic behavior in broad regions of the space of possible systems. The current hope is that analysis of broad reaches of grammar space, by sampling "random grammars", will yield deep insight into this astonishingly rich class of systems.

The promise of these random grammar systems extend from analysis of evolving proto-living systems, to characterizing mental processes such as multiple personalities, the study of technological coevolution, bounded rationality and non-equilibrium price formation at the foundations of economic theory, to cultural evolution. And the origin of life model itself, based on the probability that an arbitrary protein catalyzes an arbitrary reaction, spawned the idea of applied molecular evolution the radical concept that we might generate trillions of random genes, RNA sequences, and proteins, and learn to evolve useful polymers able to serve as drugs, vaccines, enzymes, and biosensors. The practical implications now appear large.

Strings of symbols which act upon one other to generate strings of symbols can, in general, be computationally universal. That is, such systems can carry out any specified algorithmic computation. The immense powers and yet surprising limits, trumpeted since Gödel, Turing, Church, Kleene, lie before us, but in a new and suggestive form. Strings acting on strings to generate strings create an utterly novel conceptual framework in which to cast the world. The puzzle of mathematics, of course, is that it should so often be so outrageously useful in categorizing the world. New conceptual schemes allow starkly new questions to be posed.

A grammar model is simply specified. It suffices to consider a set of M pairs of symbol strings, each about N symbols in length. The meaning of the grammar, a catch-as-catch-can set of "laws of chemistry" is this: Wherever the left member of such a pair is found in some symbol string in a "soup" of strings, substitute the right member of the pair. Thus, given an initial soup of strings, one application of the grammar might be carried out by us, acting Godlike. We regard each string in the soup in turn, try all grammar rules in some precedence order, and carry out the transformations mandated by the grammar. Strings become strings become strings. But we can let the strings themselves act on one another. Conceive of a string as an "enzyme" which acts on a second string as a "substrate" to produce a "product". A simple specification shows the idea. If a symbol sequence on a string in the soup, say 111, is identical to a symbol sequence on the "input" side of one grammar pair, then that 111 site in the string in the soup can act as an enzymatic site. If the enzymatic site finds a substrate string bearing the same site, 111, then the enzyme acts on the substrate and transforms its 111 to the symbol sequence mandated by the grammar, say 0101. Here, which symbol string in the soup acts as enzyme and which is substrate is decided at random at each encounter. With minor effort, the grammar rules can be extended to allow one enzyme string to glue two substrate strings together, or to cleave one substrate string into two product strings.

Grammar string models exhibit entirely novel classes of behavior, and all the phase transitions shown in the origin of life model. Fix a grammar. Start the soup with an initial set of strings. As these act on one another, it might be the case that all product strings are longer than all substrate strings. In this case, the system never generates a string previously generated. Call such a system a jet. Jets might be finite, the generation of strings petering out after a while, or infinite. The set of strings generated from a sustained founder set might loop back to form strings formed earlier in the process, by new pathways. Such "mushrooms" are just the autocatalytic sets proposed for the origin of life. Mushrooms might be finite or infinite, and might, if finite, squirt infinite jets into string space. A set of strings might generate only itself, floating free like an egg in string space. Such an egg is a collective identity operator in the complex parallel processing algebra of string transformations. The set of transformations collectively specifies only itself. The egg, however, might wander in string space, or squirt an infinite jet. Perturbations to an egg, by injecting a new string, might be repulsed, leaving the egg unchanged, or might unleash a transformation to another egg, a mushroom, a jet. Similarly, injection of an exogenous string into a finite mushroom might trigger a transformation to a different finite mushroom, or even an infinite mushroom. A founder set of strings might galvanize the formation of an infinite set of strings spread all over string space, yet leave local "holes" in string space because some strings might not be able to be formed from the founder set. Call such a set a filligreed fog. It may be formally undecidable whether a given string can be produced from a founder set. Finally, all possible strings might ultimately be formed, creating a pea soup in string space.

Wondrous dreamlike stuff, this. But more lies to hand. Jets, eggs, filligreed fogs and the like are merely the specification of the string contents of such an evolving system, not its dynamics. Thus, an egg might regenerate itself in a steady state, in a periodic oscillation during which the formation of each string waxes and wanes cyclically, or chaotically. The entire "edge of chaos" story concerned dynamics only, not composition. String theory opens new conceptual territory.

3. New Territory

Models of mind, models of evolution, models of technological transformation, of cultural succession, these grammar models open new provinces for precise thought. In "Origins of Order" I was able only to begin to discuss the implications of Fontana's invention. I turn next in this essay to mention their possible relation to artificial intelligence and connectionism, sketch their possible use in the philosophy of science, then discuss their use in economics, where they may provide an account, not only of technological evolution, but of bounded rationality, non-equilibrium price formation, future shock, and perhaps most deeply, a start of a theory of "individuation" of coordinated clusters of processes as entities, firms, organizations, so as to optimize wealth production. In turn, these lead to the hint of some rude analogue of the second law of thermodynamics, but here for open systems which increase order and individuation to maximize something like wealth production.

Not the least of these new territories might be a new model of mind. Two great views divide current theories of mind. In one, championed by traditional artificial intelligence, the mind carries out algorithms in which condition rules act on action rules to trigger appropriate sequences of actions. In contrast, connectionism posits neural networks whose attractors are classes, categories, or memories. The former are good at sequential logic and action, the latter are good at pattern recognition. Neither class has the strengths of the other. But parallel processing symbol strings have the strength of both. More broadly, parallel processing string systems in an open coevolving set of strings, wherein individuation of coordinated clusters of these production processes arise, may be near universal models of minds, knower and known, mutually creating the world they inhabit.

Next, some comments about the philosophy of science by an ardent amateur. Since Quine we have lived with holism in science, the realization that some claims are so central to our conceptual web that we hold them well neigh unfalsifiable, hence treat them as well neigh true by definition. Since Kuhn we have lived with paradigm revolutions and the problems of radical translation, comparability of terms before and after the revolution, reducability. Since Popper we have lived ever more uneasily with falsifiability and the injunction that there is no logic of questions. And for decades now we have lived with the thesis that conceptual evolution is like biological evolution: Better variants are cast up, never mind how conceived, and passed through the filter of scientific selection. But we have no theory of centrality versus peripherality in our web of concepts, hence no theory of pregnant versus trivial questions, nor of conceptual recastings which afford revolutions or wrinkles. But if we can begin to achieve a body of theory which accounts for both knower and known as entities which have coevolved with one another, E. coli and its world, I.B.M. and its world, and understand what it is for such a system to have a "model" of its world via meaningful materials, toxins, foods, shadow of a hawk cast on newborn chick, we must be well on our way to understanding science too as a web creating and grasping a world.

Holism should be interpretable in statistical detail. The centrality of Newton's laws of motion compared to details of geomorphology in science find their counterpart in the centrality of the automobile and peripherality of pet rocks in economic life. Conceptual revolutions are like avalanches of change in ecosystems, economic systems, and political systems. We need a theory of the structure of conceptual webs and their transformation. Pregnant questions are those which promise potential changes propagating far into the web. We know a profound question when we see one. We need a theory, or framework, to say what we know. Like Necker cubes, alternative conceptual webs are alternative grasped worlds. We need a way to categorize "alternative worlds" as if they were alternative stable collective string production systems, eggs or jets. Are mutually exclusive conceptual alternatives, like multiple personalities, literally alternative ways of being in the world. What pathways of conceptual change flow from a given conceptual web to what "neighboring" webs, and why? This is buried in the actual structure of the web at any point. Again, we know this, but need a framework to say what we know. I suspect grammar models and string theory may help. And conceptual evolution is like cultural evolution. I cannot help the image of an isolated society with a self consistent set of roles and beliefs as an egg shattered by contact with our supracritical Western civilization.

Now to economic webs, where string theory may provide tools to approach technological evolution, bounded rationality, non-equilibrium price formation, and perhaps individuation of "firms". I should stress that this work is just beginning. The suspected conclusions reach beyond that which has been demonstrated mathematically or by simulations.

A first issue is that string theory provides tools to approach the fundamental problem of technological evolution. Theoretical economists can earn a living scratching equations on blackboards. This is a strange way to catch dinner. One hundred thousand years ago, their grandfathers and grandmothers scratched a living in a more direct way. Economists survive only because the variety of goods and services in an economy has expanded since Neanderthal to include the services mathematical economists. Why? And what role does the structure of an economic web play in its own growth?

The central insight is that, in fact, the structure of an economic web at any moment plays the central role in its own transformation to a new web with new goods and services and lacking old goods and services. But it is precisely this central fact that the economists have, until now, no coherent means to think about. The richness of economic webs has increased. Introduction of the automobile, as noted, unleashes an avalanche of new goods and services ranging from gas stations to motels, and drives out horse, buggy, and the like. Economists treat technological evolution as "network externalities". This cumbersome phrase means that innovation is imagined to occur due to causes "outside" the economy. While innovation has cascading consequences of the utmost importance, traditional economic theory is not to account for technological evolution, but to note its history and treat such innovation as exogenous. Strange, since the bulk of economic growth in the current century is driven by innovation.

There is a profound reason why economics has had a difficult time building a theory of the evolution of technological webs. They lack a theory of technological complementarity and substitutability without which no such web theory can be built. String theory offers such a framework. Economists call nut and bolt, ham and eggs, "complements". That is, complements are goods and services which are used together for some purpose. Screw and nail are "substitutes", each can replace the other for most purposes. But the growth of technological niches rests on which goods and services are complements and substitutes for one another. Thus, the introduction of the computer led to software companies because software and hardware are complements. Without a theory of which goods and services are complements and substitutes for one another, one cannot build a decent account of the way technological webs grow autocatalytically.

String theory to the rescue. Any random grammar, drawn from the second order infinite set of possible grammars, can be taken not only as a catch-as-catch-can model of the "laws of chemistry", but of the unknown "laws of technological complementarity and substitutability". Strings which act on strings to make strings are tools, or capital goods. The set of strings needed as inputs to a tool to made product strings is itself a set of complements. Each string is needed with the rest to make the products. Strings which can substitute for one another as inputs to a tool to yield the same products are substitutes for one another. Such complements and substitutes constitute the "production functions" of the economist, or, for consumption goods, the consumption complementarities and substitutions, ham and eggs, salt and potassium chloride.

We have no idea what the laws of technological complementarity and substitutability are, but by scanning across grammar space we are scanning across possible models of such laws. If vast regimes of grammar space yield similar results, and we can map the regimes onto real economic systems, then those regimes of grammar space capture, in an "as if" fashion, the unknown laws of technological complementarity and substitutability which govern economic links. An ensemble theory again. Catch-as-catch-can can catch the truth.

These economic string models are now in use with Paul Romer to study the evolution of technological webs. The trick is to calculate, at each period, which of the goods currently produced, or now rendered possible by innovation based on the current goods, are produced in the next period and which current goods are no longer produced. This allows studies of avalanches of technological change.

In more detail, an economic model requires production functions, an assignment of utility to each good or service, and budget constraints. An economic equilibrium is said to exist if a ratio of the amounts of goods and services produced is found which simultaneously optimizes utility and such that all markets clear no bananas are left rotting on the dock, no hungry folk hankering for unproduced pizzas. At each period those goods, old and now possible, which are profitable are incorporated into the economy, those old and new goods which would operate at a loss will not. Thus, these models are literally the first which will show the ways avalanches of goods and services come into and leave economic systems. Heretofore, the economists have lacked a way to say why, when the automobile enters, such and such numbers of other new goods are called into existence, while old goods are rendered obsolete. Now we can study such transformations. The evolution of economic webs stands on the verge of become an integral feature of economic theory. Since such evolution dominates late 20th century and will dominate early 21st century economic growth, the capacity to study such evolution is not merely of academic interest.

String theory provides new insights into economic take off. The 21st century will undoubtedly witness the encroaching struggle between North and South, developed and underdeveloped economies, to learn how to share wealth and, more essentially, learn how to trigger adequate economic growth in the South. For decades economists have sought adequate theories of economic take off. Typically these rest on the idea of accumulation of sufficient surplus to invest. But string theory suggests this picture is powerfully inadequate. Summon all the surpluses one wishes, if the economic web is too simple in variety to allow the existing niches to call forth innovation to create novel goods and services, the economy will remain stagnant and not take off. In short, string theory suggests that an adequate complexity of goods and services is required for phase transition to take off.

These phase transitions are simply understood, and can depend upon the complexity of one economy, or the onset of trade between two insufficiently complex economies. Think of two boxes, labeled France and England. Each has a forest of types of founder strings growing within it. If, in either country, the goods and services denoted by the founder strings is too simple, then the economy will form a faltering finite jet, which soon sputters out. Few novel goods are produced at first, then none. But if the complexity of goods and services within one country is great enough, then, like an autocatalytic set at the origin of life, it will explode into an indefinitely growing proliferation of goods and services, a mushroom, filligreed fog, etc. Thus, takeoff requires sufficient technological complexity that it can feed on itself and explode. It will occur in one economy if sufficiently complex. Or onset of trade can trigger take off. Let France and England be subcritical, but begin to exchange goods. The exchange increases the total complexity, thus the growth of new opportunities, new niches for innovation, thus may catapult the coupled economies to supracritical explosion. Takeoff.

Price formation is an unsolved problem. There is no established mechanism which assures equilibrium price formation in current economic theory. We hope to show using string theory that an adequate account requires a radical transformation. Price formation is not an equilibrium phenomenon. The proper answer rests on and optimally, but boundedly, rational economic agents, who may jointly approach a price equilibrium as best as can be achieved, but typically do not reach it. Among other implications, arbitrage opportunities must typically exist.

Here is the issue. Price equilibrium is meant to be that ratio of prices for goods, denominated in money or some good, such that if all economic agents simultaneously optimize their utility functions, all markets clear. But there is a sad, if brilliant, history here. In days of old one envisioned supply and demand curves crossing. Bargaining at the bazaar between buyer and seller was to drive price to equilibrium where supply matched demand and markets cleared. Helas, ordinary folk like butter with their bread, sometimes marmalade. Unfortunately, this linkage of consumption complementarities can mean that alteration in the price of bread alters the demand for butter. In turn, theorems show that price adjustment mechanisms at the bazaar or by an auctioneer do not approach equilibrium where markets clear, but diverge away from equilibrium given any price fluctuation. Panic among the economists, if not the markets.

General equilibrium theory, the marvelous invention of Arrow and Debreu, is odd enough. Posit all possible conditional goods, bananas delivered tomorrow if it rains in Manitoba. Posit the capacity to exchange all possible such oddities, called complete markets. Posit infinitely rational economic agents with prior expectations about the future, and it can be shown that these agents, buying and selling rights to such goods at a single auction at the beginning of time, will find prices for these goods such that markets will clear as the future unfolds. Remarkable, marvelous indeed. But it was easier in the days of hunter-gatherers.

Balderdash. In the absence of complete markets the theory fails. Worse, infinite rationality is silly, we all know it. The difficulty is that if there were a "smart" knob, it always seems better to tune that knob towards high smart. The problem of bounded rationality looms as fundamental as price formation. Since Herbert Simon coined the term all economists have known this. None, it appears, has known how to solve it.

I suspect that non-equilibrium price formation and bounded rationality are linked. Using "string theory", we hope to show for an economic web whose goods and services changes over time, that even infinitely rational agents with unbounded computer time should only calculate a certain optimal number of periods into the future, a "Tc", to decide how to allocate resources at each period. Calculation yet further into the future will leave the infinitely rational agents progressively less certain about how to allocate resources in the current period. In short, even granting the economist complete knowledge on the part of agents, there is an optimal distance into the future to calculate to achieve optimal allocation of resources. Bounded rationality not only suffices, but is optimal. There is an optimal tuning of the "smart knob". If this obtains, then the same results should extend to cases with incomplete knowledge in fixed as well as evolving economies. Thank goodness. If evolution itself tunes how smart we are, perhaps we are optimally smart for our worlds.

Our ideas can be understood by comparing the task of an infinitely rational Social Planner in an economy with a fixed set of goods and services, or production technologies, from a Social Planner in evolving economic web with new goods and services becoming possible over time. In a standard economy with a fixed set of goods and services, the Social Planner can calculate precisely how he should allocate resources in the first period, or any period thereafter, by thinking infinitely far ahead. In the case of an economy with new goods and services, if the Planner thinks too far ahead he becomes progressively more confused about how he should allocate resources in the first period. He should only think optimally far ahead: He should be optimally boundedly rational.

In a standard economic model, with unchanging goods and services, each modeled as a symbol string and endowed with a utility, the Social Planner proceeds as follows: In order to allocate resources in the first period, he calculates 1 period ahead and assesses allocation of resources in the first period. Then he calculates 2 periods ahead to see how this further calculation changes the optimal allocation of resources in the first period. Then he calculates 3, 4, T periods ahead. At each such calculation he obtains an optimal ratio or allocation of economic production activities for the first period which optimizes the utility attained by action in the first period. The most important result is this: As he calculates ever further ahead, this ratio of activities at first jumps around, then settles down to a steady ratio as T approaches infinity. Two features are important. First, the further out he calculates, the larger T is, the higher the utility achieved by allocation in the first period, since he has accounted for more consequences. Second, because the ratio settles down asymptotically, that asymptotic ratio of activities at T = infinity is the optimal allocation of resources in the first period. Given this, he carries out the allocation and passes to the second period.

The central conclusion of the standard problem is that the Social Planner should tune the smart knob to maximum. A next standard step is to assume a large number of independent economic agents, each infinitely rational, each carrying out the same computation as the Social Planner. All calculate the same optimal ratio of economic activities, each does the calculated amount of his activity, utility is optimized, and because each has computed the same ratio of all activities, those activities are coordinated among the independent agents such that markets clear.

In this context, the major approach taken by economists to the fact of bounded rationality is to assume a cost of computation such that it may not be worth thinking further. The increase in utility is balanced by the cost of computing it. Such a cost is a trivial answer to why bounded rationality occurs. The deep answer, I think, is that too much calculation makes things worse. The Social Planner can be too smart by half, indeed by three quarters, or other amounts.

In an economic web where goods and services evolve over time due to innovation and replacement we hope to show that the ratio of activity calculated by the Social Planner generically will not settle down to a fixed asymptote. Rather, the further out he calculates, the more the ratio thought to be the optimal allocation of activities for the first period should jump around. Consequently, if independent economic agents carry out the same calculation as the Social Planner, the further out they calculate the harder it will become to coordinate activities. Thus, individual agents should only calculate an optimal time ahead, when the jumpiness of the optimal ratio of activities is minimized.

Here it is more slowly. As the Planner calculates for periods, T, ever further into the future in order to allocate resources in the first period, at first, as T increases, the optimal ratio will appear to settle down, but then as the hoards of new goods and services which might enter proliferate, the ratio should begin to change ever more dramatically as he calculates ever further into the future. At period 207 just the good which renders our current good 3 utterly critical makes its appearance. We would miss a gold mine had we not considered things until period 207. But upon studying period 208 we find a substitute for good 3 has become possible. We should not make 3 in the first period then. Alas again. The more we calculate the less we know.

But if the calculated ratio of activities producing goods and services first starts to settle down then becomes more variable as the Planner calculates further into the future, how in fact should he allocate resources in the first period? Every deeper calculation he becomes more confused. If he continues to calculate to infinity, even with discounting of future utilities, he may change his mind every further period he calculates.

The problem is not overwhelming for the planner, however, for he is the single commander of the entire economy, hence suffers no problems in coordinating activities across the economy. If he picks any large future T to calculate, say 1000 periods, he will make a very good allocation of resources at the current moment. Where, then, is the profound problem?

The profound problem is that there is no Social Planner. Let there be an economic agent in charge of each production function, and N such agents in the economy. Suppose, as economists do, that these agents cannot talk to one another, that they know as much as the Social Planner, and can only interact by actions. How should they coordinate their mutual behaviors? Each makes the same calculations as does the Social Planner. Each realizes that the further out he calculates, the more the ratio of activities varies. He must choose some T and act on it. But if he tries to optimize utility by choosing a large T, and others in the economy choose even slightly different T values, then each will elect to produce levels of outputs which do not mesh with the inputs assumed by others. Vast amounts of bananas will rot on the dock, hunger for apples will abound. Massive market disequilibrium occurs.

Optimally bounded rationality provides the answer. There is some period, Tc, in the future, say 7 periods ahead, when the ratio of activities is the most settled down it shall be as T varies from 1 to infinity. Here, slight differences in T chosen by other agents minimizes the bananas left on the dock and hunger for apples. Near here, then, a finite calculation ahead, is the best guess at how to allocate resources so that markets nearly clear. Bounded rationality, I believe, is linked to non-equilibrium price formation.

More should fall out. Among these, future shock. As the goods and services in the web explode in complexity, Tc should become smaller.

The time horizon for rational action would become crowded into the present.

Perhaps most importantly, we may achieve a theory of something like "optimal firm size". In a webbed economy, let local patches of production functions, vertically and horizontally integrated, count as "firms". Coordination within a firm on the choice of T can be obtained by the C.E.O. We hope that there will be an optimal distribution of firm sizes which optimizes utility. Tiny firms must remain close to Tc to minimize waste. Larger firms can push beyond Tc. But, if the economy has too few firms, hence each too large, then fluctuations over successive periods of play will drive them into bankruptcy. Thus, an intermediate number of firms, of intermediate size, should optimize average self and mutual wealth production over time.

But a theory of firm size as a cluster of production processes which optimizes the distribution of firm sizes such that each "patch" optimally maximizes growth of utility is no small conceptual step. It is a start toward a theory of individuation of clustered sets of production processes as "entities " which optimally coevolve with one another in the economic system. As such, it seems deeply linked to coevolution to the edge of chaos. In both cases, tuning something like the surface to volume ratio of an "individual" such that all individuated entities optimize expected success, is the key . More, such a theory of individuation hints an analogue of the second law of thermodynamics for open thermodynamic systems. Such a law is one ultimate focus of the "sciences of complexity".

The root requirement is the primitive concept of "success", taken as optimizing utility in economics, or optimizing reproduction success in biology. Let red and blue bacteria compete, while the former reproduces more rapidly. Soon the Petri plate is red. Let the bacteria not divide but increase in mass. Again, soon the Petri plate is red. Increase of mass is the analogue of increase of wealth. It is not an accident that biology and economics borrow from one another. The wealth of nations is, at base, the analogue of the wealth of a species. More is better.

Given a primitive concept of success, then a theory of parallel processing algorithms, alive in the space of possible transformations, should yield a theory of individuals as clumps or patches of processes which coordinate behavior such that each optimizes success, while all optimize success as best they can in the quivering shimmering world they mutually create. Economics is the purest case, for coordination into firms is a voluntary process. But Leo Buss has stressed the puzzle of the evolution of multicellular individuals since the Cambrian. Why should an individual cell, capable of dividing and passing its genome towards the omega point of time, choose to forego that march and enter into a multicellular organism where it shall form somatic tissue, not gonadal tissue, hence die. Think it not so? The slime mold Dicteostylium discoidum coalesces thousands of starving amoebae, each capable of indefinite mitotic division, into a crawling slug in which many cells later form stalk, not spore, hence die. Their progeny are cut off. Yet they have opted to associate into a biological firm, replete with specialization of labor and termination with extreme disfavor, for some form of profit. It is, at base, the same problem. What sets the size, volume, and boundary membrane of an individual.

Given a theory of individuals, patches of coordinate processes optimizing success in a coevolving world mutually known, then the argument above and its generalizations in the case of incomplete knowledge and error amplification with excessive calculation, yield a bound to rationality. A coevolving individual does not benefit, nay, does worse, by calculating too far into the future. Or too far into the web away from each. Or, equally, too far into the future light cone of events. But, in turn, a bound on rationality, better, an optimally bounded rationality, implies a bound on the complexity of the coevolving individual. No point in being overcomplex relative to one's world. The internal portrait, condensed image, of the external world carried by the individual and used to guide its interactions, must be tuned, just so, to the ever evolving complexity of the world it helped create.

We draw near an analogue of the second law of thermodynamics. The latter fundamental law states that closed systems approach a state of disorder, entropy increases to a maximum. Living systems, E.coli or IBM, are open systems. Matter and energy range through each as the precondition for their emergence. The investigations sketched here intimate a law of increasing order and differentiation of "individuals", packages of processes, probably attaining the boundary of chaos, a wavefront of self organizing processes in the space of processes, molecular, economic, cultural, a wavefront of lawful statistical form governed by the generalized insight of Darwin and Smith. While the attainment of optimally sized "individual" which optimize coevolution may be constrained by the means allowing individuals to form, aggregation of cells, antitrust laws, the direction of optimization, like the direction of entropy change, will govern the emerging structure.

4. Closing Remark: A Place for Laws in Historical Sciences

I close this essay by commenting on Burian and Richardsons' thoughtful review of *Origins of Order*. They properly stress a major problem: What is specifically "biological" in the heralded renderings of ensemble theories? This is a profound issue. Let me approach it by analogy with the hoped for use of random grammar models in economics as discussed above. As emphasized, economists lack a theory of technological evolution because they lack a theory of technological complementarities and substitutes. One needs to know why nuts go with bolts to account for the coevolution of these two bits of econo-stuff. But we have no such theory, nor is it even clear what a theory which gives the actual couplings among ham and eggs, nuts and bolts, screws and nails, computer and software engineer, might be. The hope for grammar models is that each grammar model, one of a non-denumerably infinite set of such grammars since grammars map power sets of strings into power sets of strings, each such grammar model is a "catch-as-catch-can" model of the unknown laws of technological complementarity and substitutability. The hope is that vast reaches of "grammar space" will yield economic models with much the same global behavior. If such generic behaviors map onto the real economic world, I would argue that we have found the proper structure of complementarity and substitutability relationships among goods and services, hence can account for many statistical aspects of economic growth. But this will afford no account of the coupling between specific economic goods such as power transmission and the advent of specific new suppliers to Detroit. Is such a theory specifically "economics"? I do not know, but I think so.

Grammar models afford us the opportunity to capture statistical features of deeply historically contingent phenomena ranging from biology to economics, perhaps to cultural evolution. Phase transitions in complex systems may be lawful, power law distributions of avalanches may be lawful, but the specific avalanches of change may not be predictable. Too many throws of the quantum dice. Thus we confront a new conceptual tool which may provide a new way of looking for laws in historical sciences. Where will the specifics lie? As always, I presume, in the consequences deduced after the axioms are interpreted.

References

Bagley, R. (1991), *A Model of Functional Self Organization.* Ph.D Thesis, University of California, San Diego.

Dawkins, R. (1976), *The Selfish Gene.* Oxford University Press, Oxford, N.Y.

Derrida, B. and Pomeau, Y. (1986), "Random networks of automata: a simple annealed approximation." *Europhys. Letters.* 1(2): 45-49.

Edwards, D.F., and Anderson, P.W. (1975), *Journal of Physics* F 5:965.

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

Fogleman-Soulie, F. (1985), "Parallel and sequential computaiton in Boolean networks." In *Theoretical Computer Science 40*, North Holland.

Fontana, W. (1991), "Artificial Life II", in Langton, Farmer, Taylor (eds.) in press. Addison Wesley.

Hopfield, J.J. (1982), *Proceedings of the National Academy of Science.* U.S.A. 79: 2554-2558.

Jacob, F., and Monod, J. (1961), "On the regulation of gene activity." Cold Spring Harbor Symposium. *Quantum Blology* 26: 193-211.

______________. (1963), "Genetic repression, allosteric inhibition, and cellular differentiation." In È. (M. Locke, ed.) 21st Symposium for the Society for the Study of Development and Growth. Academic Press, N.Y. pp. 30-64.

Kauffman, S.A. (1969), "Metabolic stability and epigenesis in randomly connected nets." *Journal of Theoretical Biology* 22: 437-467.

_________. (1986) "Autocatalytic sets of Protiens." *Journal Theoretical Biology* 119: 1-24.

_________. (1989), "Principles of Adaptation in Complex Systems." In *Lectures in the Sciences of Complexity* in Dan Stein (ed.) The Santa Fe Institute Series. Addison Wesley.

_________. (1991), "Antichaos and Adaptation." *Scientific American*, August 1991.

_________. (1992), *Origins of Order: Self Organization and Selection in Evolution*. In press, Oxford University Press.

Langton, C. (1991), in *Artificial Life II*, (eds.) Langton, Farmer, Taylor, in press, Addison Wesley.

Smolensky, P. (1988), "On the proper treatment of connectionism." *Behavioral and Brain Science* 11: 1-74.

Stauffer, D. (1987), "Random Boolean networks: analogy with percolation." *Philosophical Magazine* B, 56 no. 6: 901-916.