

# Complexity—An Introduction

John L Casti

## I. COMPLEXITY AS A SYSTEMS CONCEPT

In everyday parlance the term “complex” is generally taken to mean a person or thing composed of many interacting components whose behavior and/or structure is difficult to understand. The behavior of national economies, the human brain and a rain forest ecosystem are all good illustrations of complex systems. These examples underscore the point that sometimes a system may be structurally complex, like a mechanical clock, but behave very simply. In fact, it’s the simple, regular behavior of a clock that allows it to serve as a timekeeping device. On the other hand, there are systems, like the toy rotator shown in Figure 1, whose structure is very easy to understand but whose behavior is impossible to predict. And, of course, some systems like the brain are complex in both structure and behavior.

The examples just cited show that there’s nothing new about complex systems; they’ve been with us from the time our ancestors crawled up out of the sea. But what is new is that for perhaps the first time in history, we have the knowledge—and the tools—to study such systems in a controlled, repeatable, scientific fashion. So there is reason to believe that this newfound capability will eventually lead to a viable theory of such systems.

Prior to the recent arrival of cheap and powerful computing capabilities, we were hampered in our ability to study a complex system like a road-traffic network, a national




Figure 1. A toy rotator with unpredictable behavior.

economy or a supermarket chain because it was simply too expensive, impractical, too time-consuming—or too dangerous—to tinker with the system as a whole. Instead, we were limited to biting off bits and pieces of such processes that could be looked at in a laboratory or in some other controlled setting. But with today’s computers we can actually build complete silicon surrogates of these systems, and use these “would-be worlds” as laboratories within which to look at the workings—and behaviors—of the complex systems of everyday life.

In coming to terms with complexity as a systems concept, we first have to realize that complexity is an inherently subjective concept; what’s complex depends upon how you look. When we speak of something being complex, what we’re really doing is making use of everyday language to express a feeling or impression that we dignify with the label “complex.” But the meaning of something depends not only on the language in which it is expressed (i.e., the code), the medium of transmission and the message, but also on the context. In short, meaning is bound up with the whole process of communication and doesn’t reside in just one or another aspect of it. As a result, the complexity of a political

structure, an ecosystem or an immune system cannot be regarded as simply a property of that system taken in isolation. Rather, whatever complexity such systems have is a joint property of the system *and* its interaction with another system, most often an observer and/or controller.

This point is easy to see in areas like finance. An individual investor interacts with the stock exchange and thereby affects the price of a stock by deciding to buy, to sell or to hold. This investor then sees the market as complex or simple, depending on how he or she perceives the change of prices. But the exchange itself acts upon the investor, too, in the sense that what is happening on the floor of the exchange influences the investor's decisions. This *back interaction* causes the market to see the investor as having a certain degree of complexity, in that the investor's actions cause the market to be described in terms like *nervous*, *calm* or *unsettled*. The kind of two-way complexity of a financial market becomes especially obvious in situations when the investor is one whose trades make noticeable blips on the ticker without actually dominating the market.

So just as with truth, beauty, good and evil, complexity resides as much in the eye of the beholder as it does in the structure and behavior of a system itself. This is not to say that there do not exist *objective* ways to characterize some aspects of a system's complexity. After all, an amoeba is just plain simpler than an elephant by whatever notion of complexity you happen to believe in. The main point, though, is that these objective measures only arise as special cases of the two-way measures, cases in which the interaction between the system and the observer is much weaker in one direction than in the other.

A second key point is that common usage of the term *complex* is informal. The word is typically employed as a name for something that seems counterintuitive, unpredictable

or just plain hard to understand. So if it's a genuine *science* of complex systems we're after and not just anecdotal accounts based on vague personal opinions, we're going to have to translate some of these informal notions about the complex and the commonplace into a more formal, stylized language, one in which intuition and meaning can be more or less faithfully captured in symbols and syntax. The problem is that an integral part of transforming complexity (or anything else) into a science involves making that which is fuzzy precise, not the other way around, an exercise we might more compactly express as "formalizing the informal."

To bring home this point a bit more forcefully, let's consider some of the properties associated with *simple* systems by way of inching our way to a feeling for what's involved with the complex. Generally speaking, simple systems exhibit the following characteristics:

- *Predictable behavior*—There are no surprises in simple systems; simple systems give rise to behaviors that are easy to deduce if we know the inputs (decisions) acting upon the system and the environment. If we drop a stone, it falls; if we stretch a spring and let it go, it oscillates in a fixed pattern; if we put money into a fixed-interest bank account, it grows to a predictable sum in accordance with an easily understood and computable rule. Such predictable and intuitively well-understood behavior is one of the principal characteristics of simple systems.

Complex processes, on the other hand, generate counterintuitive, seemingly acausal behavior that's full of surprises. Lower taxes and interest rates lead to higher unemployment; low-cost housing projects give rise to slums worse than those the "better" housing

replaced; the construction of new freeways results in unprecedented traffic jams and increased commuting times. For many people, such unpredictable, seemingly capricious, behavior is the defining feature of a complex system.

- *Few interactions and feedback/feedforward loops*—Simple systems generally involve a small number of components, with self-interactions dominating the linkages among the variables. For example, primitive barter economies, in which only a small number of goods (food, tools, weapons, clothing) are traded, seem much simpler and easier to understand than the developed economies of industrialized nations, in which the pathways between raw material inputs and finished consumer goods follow labyrinthine routes involving large numbers of interactions between various intermediate products, labor and capital inputs.

In addition to having only a few variables, simple systems generally consist of very few feedback/feedforward loops. Loops of this sort enable the system to restructure, or at least modify, the interaction pattern among its variables, thereby opening up the possibility for a wider range of behaviors. To illustrate, consider a large organization that's characterized by variables like employment stability, substitution of capital for human labor and level of individual action and responsibility (individuality). Increased substitution of work by capital decreases the individuality in the organization, which in turn may reduce employment stability. Such a feedback loop exacerbates any internal stresses initially present in the system, leading possibly to a collapse of the entire organization. This type of collapsing loop is especially dangerous for social structures, as it threatens their ability to absorb shocks, which seems to be a common feature of complex social phenomena.

- *Centralized decisionmaking*—In simple systems, power is generally concentrated in one or at most a few decisionmakers. Political dictatorships, privately owned corporations and the Roman Catholic Church are good examples of this sort of system. These systems are simple because there is very little interaction, if any, between the lines of command. Moreover, the effect of the central authority's decision upon the system is usually rather easy to trace.

By way of contrast, complex systems exhibit a diffusion of real authority. Such systems seem to have a nominal supreme decision-maker, but in actuality the power is spread over a decentralized structure. The actions of a number of units then combine to generate the actual system behavior. Typical examples of these kinds of systems include democratic governments, labor unions and universities. Such systems tend to be somewhat more resilient and stable than centralized structures because they are more forgiving of mistakes by any one decision-maker and are more able to absorb unexpected environmental fluctuations.

- *Decomposable*—Typically, a simple system involves weak interactions among its various components. So if we sever some of these connections, the system behaves more or less as before. Relocating American Indians to reservations produced no major effects on the dominant social structure in New Mexico and Arizona, for example, since, for various cultural reasons, the Indians were only weakly coupled to the dominant local social fabric in the first place. Thus the simple social interaction pattern present could be further decomposed and studied as two independent processes—the Indians and the settlers.

Complex processes, on the other hand, are irreducible. Neglecting any part of the process or severing any of the connections linking its parts usually destroys essential aspects of the system’s behavior or structure. The N-Body Problem in physics is a quintessential example of this sort of indecomposability. Other examples include an electrical circuit, a Renoir painting or the tripartite division of the U.S. government into its executive, judicial and legislative subsystems. You just can’t start slicing up systems of this type into subsystems without suffering an irretrievable loss of the very information that makes these systems a “system.”

## II. SURPRISE-GENERATING MECHANISMS

The vast majority of counterintuitive behaviors shown by complex systems are attributable to some combination of the following five sources: paradoxes/self-reference, instability, uncomputability, connectivity, and emergence. With some justification, we can think of these sources of complexity as *surprise-generating mechanisms*, whose quite different natures each lead to their own characteristic type of surprise. Let’s take a quick look at each of these mechanisms before turning to a more detailed consideration of how they act to create complex behavior.

*Paradox:* Paradoxes arise from false assumptions about a system leading to inconsistencies between its observed behavior and our *expectations* of that behavior. Sometimes these situations occur in simple logical or linguistic situations, such as the famous Liar Paradox (“This sentence is false.”). In other situations, the paradox comes from the peculiarities of the human visual system, as with the impossible staircase shown in Figure 2, or




Figure 2. The impossible staircase.

simply from the way in which the parts of a system are put together, like the developing economy discussed in the preceding section.

*Instability:* Everyday intuition has generally been honed on systems whose behavior is stable with regard to small disturbances, for the obvious reason that unstable systems tend not to survive long enough for us to develop good intuitions about them. Nevertheless, the systems of both nature and humans often display pathologically sensitive behavior to small disturbances, as for example, when stock markets crash in response to seemingly minor economic news about interest rates, corporate mergers, or bank failures. Such behaviors occur often enough that they deserve a starring role in our taxonomy of surprise. Here is a simple example illustrating the point.

In Adam Smith's (1723–1790) world of economic processes, a world involving a system of goods and demand for those goods, prices will always tend toward a level at which supply equals demand. Thus, this world postulates some type of negative feedback from the supply/demand relationship to prices, which leads to a level of prices that is stable. This means that any change in prices away from this equilibrium will be resisted by the



economy, and that the laws of supply and demand will act to reestablish the equilibrium prices. Recently, economists have argued that this is not the way many of the sectors work in the real economy at all. Rather, these economists claim that what we see is *positive* feedback in which the price equilibria are unstable.

For example, when video cassette recorders (VCRs) started becoming a household item some years back, the market began with two competing formats—VHS and Beta—selling at about the same price. By increasing its market share, each of these formats could obtain increasing returns since, for example, large numbers of VHS recorders would encourage video stores to stock more prerecorded tapes in VHS format. This, in turn, would enhance the value of owning a VHS machine, leading more people to buy machines of that format. By this mechanism a small gain in market share could greatly amplify the competitive position of VHS recorders, thus helping that format to further increase its share of the market. This is the characterizing feature of positive feedback—small changes are amplified instead of dying out.

The feature of the VCR market that led to the situation described above is that it was initially unstable. Both VHS and Beta systems were introduced at about the same time and began with approximately equal market shares. The fluctuations of those shares early on were due principally to things like luck and corporate maneuvering. In a positive-feedback environment, these seemingly chance factors eventually tilted the market toward the VHS format until the VHS acquired enough of an advantage to take over virtually the entire market. But it would have been impossible to predict at the outset which of the two systems would ultimately win out. The two systems represented a pair of unstable equilibrium points in competition, so that unpredictable chance factors ended up shifting

the balance in favor of VHS. In fact, if the common claim that the Beta format was technically superior holds any water, then the market's choice did not even reflect the best outcome from an economic point of view.

*Uncomputability:* The kinds of behaviors seen in models of complex systems are the end result of following a set of rules. This is because these models are embodied in computer programs, which in turn are necessarily just a set of rules telling the machine what bits in its memory array to turn on or off at any given stage of the calculation. By definition, this means that any behavior seen in such worlds is the outcome of following the rules encoded in the program. Although computing machines are *de facto* rule-following devices, there is no *a priori* reason to believe that any of the processes of nature and humans are necessarily rule-based. If uncomputable processes do exist in nature—for example, the breaking of waves on a beach or the movement of air masses in the atmosphere—then we could never see these processes manifest themselves in the surrogate worlds of their models. We may well see processes that are close approximations to these uncomputable ones, just as we can approximate an irrational number as closely as we wish by a rational number. However, we will never see the real thing in our computers, if indeed such uncomputable quantities exist outside the pristine world of mathematics.

To illustrate what is at issue here, the problem of whether the cognitive powers of the human mind can be duplicated by a computing machine revolves about just this question. If our cognitive activity is nothing more than following rules encoded somehow into our neural circuitry, then there is no logical obstacle to constructing a “silicon mind.” On the other hand, it has been forcefully argued by some that cognition involves activities

that transcend simple rule following. If so, then the workings of the brain can never be captured in a computer program and the technophobes of the world can all rest easier at night. This is because there can then be no set of rules—a computer program—that could faithfully capture *all* the things happening in the cognitive processes of the human brain.

*Connectivity:* What makes a system a system and not simply a collection of elements is the connections and interactions among the individual components of the system, as well as the effect these linkages have on the behavior of the components. For example, it is the interrelationship between capital and labor that makes an economy. Each component taken separately would not suffice. The two must interact for economic activity to take place. As we saw with the example of the developing world economy, complexity and surprise often resides in these connections. Here is another illustration of this point.

Certainly the most famous question of classical celestial mechanics is the *N-Body Problem*, which comes in many forms. One version involves  $N$  point masses moving in accordance with Newton’s laws of gravitational attraction, and asks if from some set of initial positions and velocities of the particles, there is a finite time in the future at which either two (or more) bodies collide or one (or more) bodies acquires an arbitrarily high energy, and thus flies off to infinity. In the special case when  $N = 10$ , this is a mathematical formulation of the question, “Is our solar system stable?”

The behavior of two planetary bodies orbiting each other can be written down completely in terms of the elementary functions of mathematics, like powers, roots, sines, cosines, and exponentials. Nevertheless, it turns out to be impossible to combine the solutions of three two-body problems to determine whether a three-body system is stable.

Thus, the essence of the Three-Body Problem resides somehow in the way in which *all three* bodies interact. Any approach to the problem that severs even one of the linkages among the bodies destroys the very nature of the problem. Here is a case in which complicated behavior arises as a result of the interactions between relatively simple subsystems.

Incidentally, in a 1988 doctoral dissertation based upon earlier work by Don Saari, Jeff Xia of Northwestern University gave a definitive answer to the general question of the stability of such systems by constructing a five-body system for which one of the bodies does indeed acquire an arbitrarily large velocity after a finite amount of time. This result serves as a counterexample to the idea that perhaps *all*  $n$ -body systems are actually stable. The general idea underlying Xia's construction is shown in Figure 3, where we see two binary systems and a fifth body that shuttles back and forth between them. By arranging things just right, the single body can be made to move faster and faster between the binaries until at some finite time its velocity exceeds any predefined level.




Figure 3. Xia's solution of the N-Body Problem.

This result says nothing about the specific case of our solar system, which is certainly not configured like Xia's example. In fact, computer simulations of the formation of planetary systems suggest that most real planetary systems are very different than the special configuration used by Xia—a good example of the difference between the way a physicist and a mathematician look at the same problem! But Xia's example does suggest that *perhaps* the solar system is not stable, and more importantly offers new tools with which to investigate the matter further.

*Emergence:* A surprise-generating mechanism dependent on connectivity for its very existence is the phenomenon of emergence. This refers to the way the interactions among system components generates unexpected global system properties not present in any of the subsystems taken individually. A good example is water, whose distinguishing characteristics are its natural form as a liquid and its nonflammability, both of which are totally different than the properties of its component gases, hydrogen and oxygen.

The difference between complexity arising from emergence and that coming only from connection patterns lies in the nature of the interactions among the various component pieces of the system. For emergence, attention is not simply on whether there is some kind of interaction between the components, but also on the specific nature of that interaction. For instance, connectivity alone would not enable one to distinguish between ordinary tap water involving an interaction between hydrogen and oxygen molecules and heavy water (deuterium), which involves interaction between the same components albeit with an extra neutron thrown in to the mix. Emergence would make this distinction. In practice it is often difficult (and unnecessary) to differentiate between connectivity and emergence,

and they are frequently treated as synonymous surprise-generating procedures. A good example of emergence in action is the organizational structure of an ant colony.

Like human societies, ant colonies achieve things that no individual ant could accomplish on its own. Nests are erected and maintained, chambers and tunnels are excavated, and territories are defended. All these activities are carried on by individual ants acting in accord with simple, local information; there is no master ant overseeing the entire colony and broadcasting instructions to the individual workers. Somehow each individual ant processes the partial information available to it in order to decide which of the many possible functional roles it should play in the colony.

Recent work on harvester ants has shed considerable light on the process by which an ant colony assesses its current needs and assigns a certain number of members to perform a given task. These studies identify four distinct tasks an adult harvester-ant worker can perform outside the nest: foraging, patrolling, nest maintenance, and midden work (building and sorting the colony's refuse pile). So it is these different tasks that define the components of the system we call an ant colony, and it is the interaction among ants performing these tasks that gives rise to emergent phenomena in the colony.

One of the most notable interactions is between forager ants and maintenance workers. When nest-maintenance work is increased by piling some toothpicks near the opening of the nest, the number of foragers decreased. Apparently, under these environmental conditions the ants engaged in task switching, with the local decision made by each individual ant determining much of the coordinated behavior of the entire colony. Task allocation depends on two kinds of decisions made by individual ants. First, there is the decision about which task to perform, followed by the decision of whether to be active in this task. As already

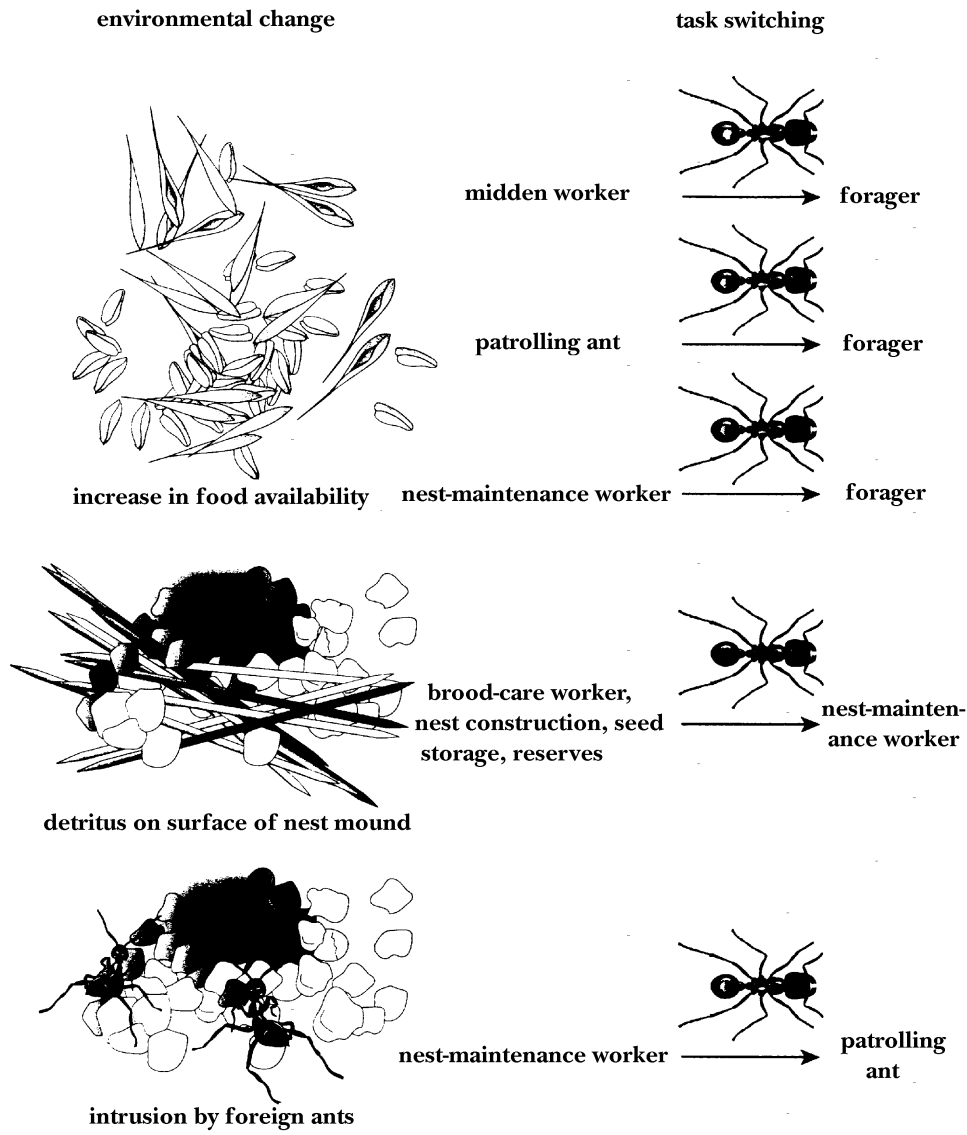


Figure 4. Task switching in a harvester ant colony.

noted, these decisions are based solely on local information; there is no central decision maker keeping track of the big picture.

Figure 4 gives a summary of the task-switching roles in the harvester ant colony, showing that once an ant becomes a forager it never switches back to other tasks outside the nest. When a large cleaning chore arises on the surface of the nest, new nest-maintenance workers are recruited from ants working inside the nest, not from workers performing tasks on the

outside. When there is a disturbance like an intrusion by foreign ants, nest-maintenance workers will switch tasks to become patrollers. Finally, once an ant is allocated a task outside the nest, it never returns to chores on the inside.

The ant colony example shows how interactions among the various types of ants can give rise to patterns of global work allocation in the colony, patterns that could not be predicted or that could not even arise in any single ant. These patterns are emergent phenomena due solely to the types of interactions among the different tasks.

Table 1 gives a summary of the surprise-generating mechanisms just outlined.

Table 1. The main surprise-generating mechanisms.

<b>Mechanism</b>	<b>Surprise Effect</b>
Paradoxes	Inconsistent phenomena
Instability	Large effects from small changes
Uncomputability	Behavior transcends rules
Connectivity	Behavior cannot be decomposed into parts
Emergence	Self-organizing patterns

### III. EMERGENT PHENOMENA

Complex systems produce surprising behavior; in fact, they produce behavioral patterns and properties that just cannot be predicted from knowledge of their parts taken in isolation. These so-called “emergent properties” are probably the single most distinguishing feature of complex systems. An example of this phenomenon occurs when one considers a collection of independent random quantities, such as the heights of all the people in New York City. Even though the individual numbers in this set are highly variable, the distribution of this set of numbers will form the familiar bell-shaped curve of elementary



statistics. This characteristic bell-shaped structure can be thought of as “emerging” from the interaction of the component elements. Not a single one of the individual heights can correspond to the normal probability distribution, since such a distribution implies a population. Yet when they are all put into interaction by adding and forming their average, the Central Limit Theorem of probability theory tells us that this average and the dispersion around it must obey the bell-shaped distribution. Another example of emergent phenomena occurs in a checkerboard system called *Life*.

In the October 1970 issue of *Scientific American* magazine, columnist Martin Gardner introduced the world to a simple board game that has come to be called the *Game of Life*. In actuality, *Life* is not really a game, since there are no players; nor are any decisions to be made. Rather it is a dynamical system played on the squares of an infinite checkerboard, technically what’s termed a 2-state, two-dimensional cellular automaton.

The rule determining whether each square on the board is ON or OFF at any given moment for this system was laid down by the game’s inventor, Cambridge University mathematician John Horton Conway. It specifies the following fates for a given cell in terms of the state of each of its 8 adjacent neighbors:

1. The cell will be ON in the next generation if exactly 3 of its neighboring cells are currently ON,
2. The cell will retain its current state if exactly 2 of its neighbors are ON,
3. The cell will be OFF otherwise.

This rule is Conway’s attempt to balance out a cell’s dying of “isolation” if it has too few living neighbors and dying of “overcrowding” if it has too many.

Figure 5 shows the fate of some initial triplet patterns using this rule. We see that the first three triplets die out after the second generation, whereas the fourth triplet forms a stable Block, and the fifth, termed a Blinker, oscillates indefinitely. Figure 6 shows the life history of the first three generations of an initially more-or-less randomly populated life universe.




Figure 5. Some triplet histories in *Life*

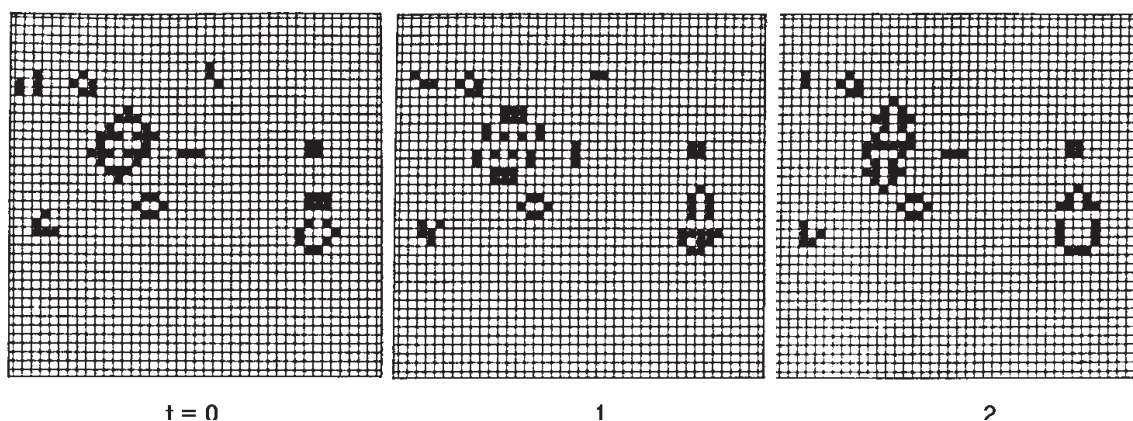


Figure 6. A typical *Life* history

Of considerable interest is whether there are initial configurations that are eventually copied elsewhere in the array. The first example of such a configuration is the Glider, which is displayed in Figure 7.




Figure 7. The Glider

In the early days of *Life*, researchers conjectured that because of the overpopulation constraint built into the Life rule, there were no configurations that could grow indefinitely. Conway offered a \$50 reward to anyone who could produce such a configuration. The configuration that won the prize, termed a Glider Gun, is depicted in Figure 8. Here the Gun, shown at the lower left of the figure, is a spatially fixed oscillator that repeats its original shape after thirty generations. Within this period, the Gun emits a Glider, which wanders across the grid and encounters an Eater, which is shown at the top right of the figure. The Eater, a fifteen generation oscillator, swallows up the Glider without undergoing any irreversible change itself. Since the Gun oscillates indefinitely, it can produce an infinite number of Gliders, implying that configurations that can grow indefinitely do exist.

The *Life* rule has led to a veritable cornucopia of strange and captivating patterns, of which the Glider Gun is just one. And each of these patterns is an example of an emergent property, one that comes from the interaction of the various squares and cannot be determined just from looking at what a single square is doing in isolation.




Figure 8. The Glider Gun

Life is an example of how emergent properties arise in a purely mathematical world, the infinite checkerboard upon which the game is played. A more real-world example of an emergent property is the movement of prices on a financial market.

Around 1988, Brian Arthur, an economist from Stanford, and John Holland, a computer scientist from the University of Michigan, hit upon the idea of creating an artificial stock market inside a computer, one that could be used to answer a number of questions that people in finance had wondered and worried about for decades. Among these questions are:

- Does the average price of a stock settle down to its so-called *fundamental value*, the value determined by the discounted stream of dividends that one can expect to receive by holding the stock indefinitely?
- Is it possible to concoct technical trading schemes that systematically turn a profit greater than a simple buy-and-hold strategy?
- Does the market eventually settle into a fixed pattern of buying and selling? In other words, does it reach “stationarity”?

Arthur and Holland knew that the conventional wisdom of finance argued that today’s price of a stock was simply the discounted *expectation* of tomorrow’s price plus dividend, given the information available about the stock today. This theoretical price-setting procedure is based on the assumption that there is an objective way to use today’s information to form this expectation. But this information typically consists of past prices, trading volumes, economic indicators and the like. So there may be *many* perfectly defensible ways based on many different assumptions to statistically process this information in order to forecast tomorrow’s price.

The simple observation that there is no single, best way to process information led Arthur and Holland to the conclusion that deductive methods for forecasting prices are, at best, an academic fiction. As soon as you admit the possibility that not all traders in the market arrive at their forecasts in the same way, the deductive approach of classical finance theory begins to break down. So a trader must make assumptions about how other investors form expectations and how they behave. He or she must try to “psyche out” the market. But this leads to a world of *subjective* beliefs and to beliefs about those beliefs. In short, it leads to a world of induction rather than deduction.

In order to answer the questions above, Arthur and Holland, along with physicist Richard Palmer, finance theorist Blake LeBaron and market trader Paul Tayler built an electronic market where they could, in effect, play god by manipulating trader’s strategies, market parameters and all the other experiments that cannot be done on real stock exchanges.

This surrogate market consists of

- a. a fixed amount of stock in a single company;
- b. a number of “traders” (computer programs) that can trade shares of this stock at each time period;
- c. a “specialist” who sets the stock price endogenously by observing market supply and demand and matching orders to buy and to sell;
- d. an outside investment (“bonds”) in which traders can place money at a varying rate of interest;

e. a dividend stream for the stock that follows a random pattern.

As for the traders, the model assumes that they each summarize recent market activity by a collection of “descriptors,” which involve verbal characterization like “the market has gone up every day for the past week,” or “the market is nervous,” or “the market is lethargic today.” Let’s label these descriptors A, B, C, and so on. In terms of the descriptors, the traders decide whether to buy or sell by rules of the form: “If the market fulfills conditions A, B, and C, then BUY, but if conditions D, G, S, and K are fulfilled, then HOLD.” Each trader has a collection of such rules, and acts on only one rule at any given time period. This rule is the one that the trader views as his or her currently most accurate rule.

As buying and selling goes on in the market, the traders can re-evaluate their different rules in two different ways: (1) by assigning higher probability of triggering a given rule that has proved profitable in the past, and/or (2) by recombining successful rules to form new ones that can then be tested in the market. This latter process is carried out by use of what’s called a “genetic algorithm,” which mimics the way nature combines the genetic pattern of males and females of a species to form a new genome that is a combination of those from the two parents.

A run of such a simulation involves initially assigning sets of predictors to the traders at random, and then beginning the simulation with a particular history of stock prices, interest rates and dividends. The traders then randomly choose one of their rules and use it to start the buying-and-selling process. As a result of what happens on the first round of trading, the traders modify their estimate of the “goodness” of their collection of rules, generate new rules (possibly) and then choose the best rule for the next round of trading.

And so the process goes, period-after-period, buying, selling, placing money in bonds, modifying and generating rules, estimating how good the rules are and, in general, acting in the same way that traders act in real financial markets. The overall flow of activity in this market is shown in Figure 9.

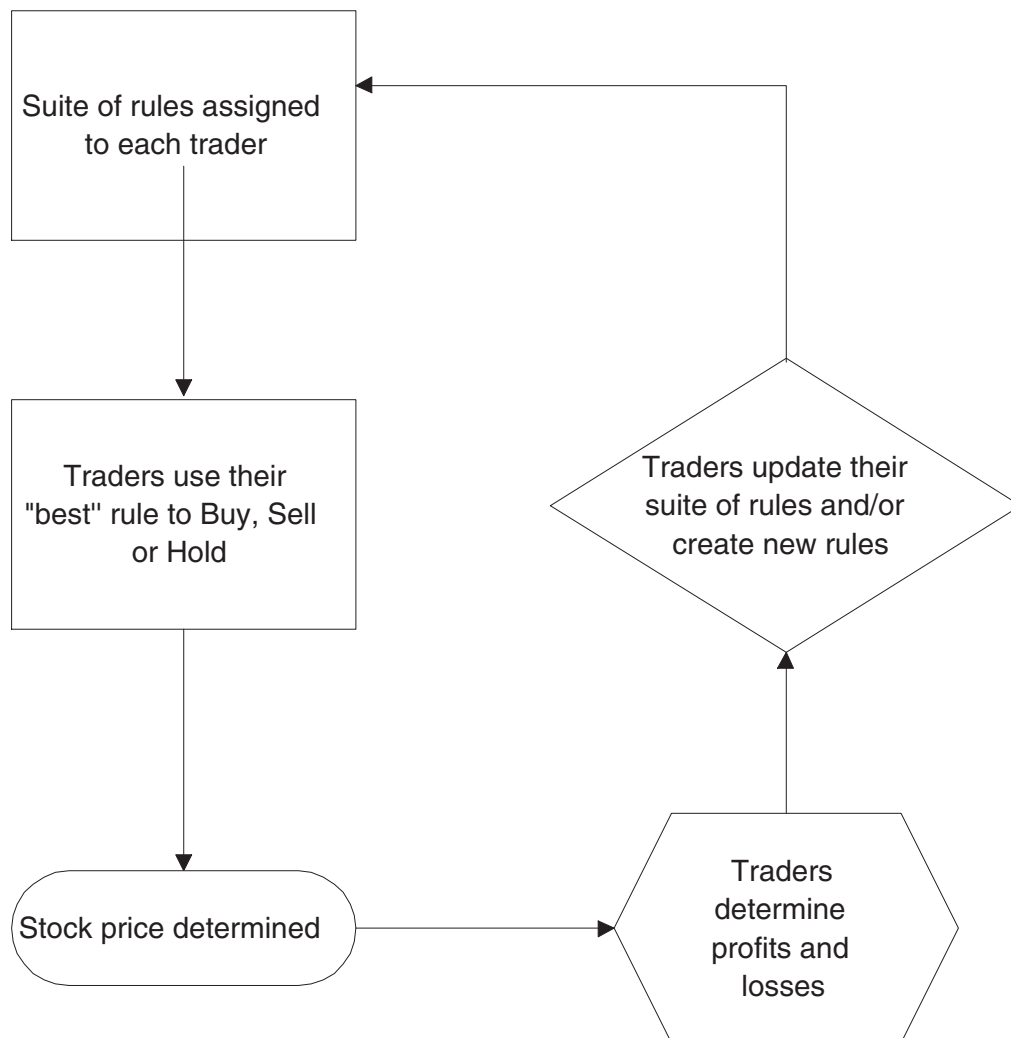


Figure 9. The logical flow of activity in the stock market.

A typical moment in this artificial market is displayed in Figure 10. Moving clockwise from the upper left, the first window shows the time history of the stock price and dividend, where the current price of the stock is the black line and the top of the grey region is the



current fundamental value. So the region where the black line is much greater than the height of the grey region represents a price bubble, while the market has “crashed” in the region where the black line sinks far below the grey. The upper right window is the current relative wealth of the various traders, while the lower right window displays their current level of stock holdings. The lower left window shows the trading volume, where grey is the selling volume and black is the buying volume. The total number of trades possible is then the minimum of these two quantities, since for every share purchased there must be one share available for sale. The various buttons on the screen are for parameters of the market that can be set by the experimenter.

After many time periods of trading and modification of the traders’ decision rules, what emerges is a kind of “ecology” of predictors, with different traders employing different rules to make their decisions. Furthermore, it is observed that the stock price always settles down to a random fluctuation about its fundamental value. But within these fluctuations a very rich behavior is seen: price bubbles and crashes, psychological market “moods,” overreactions to price movements and all the other things associated with speculative markets in the real world.

Also as in real markets, the population of predictors in the artificial market continually coevolves, showing no evidence of settling down to a single best predictor for all occasions. Rather, the optimal way to proceed at any time is seen to depend critically upon what everyone else is doing at that time. In addition, we see mutually-reinforcing trend-following or technical-analysis-like rules appearing in the predictor population.

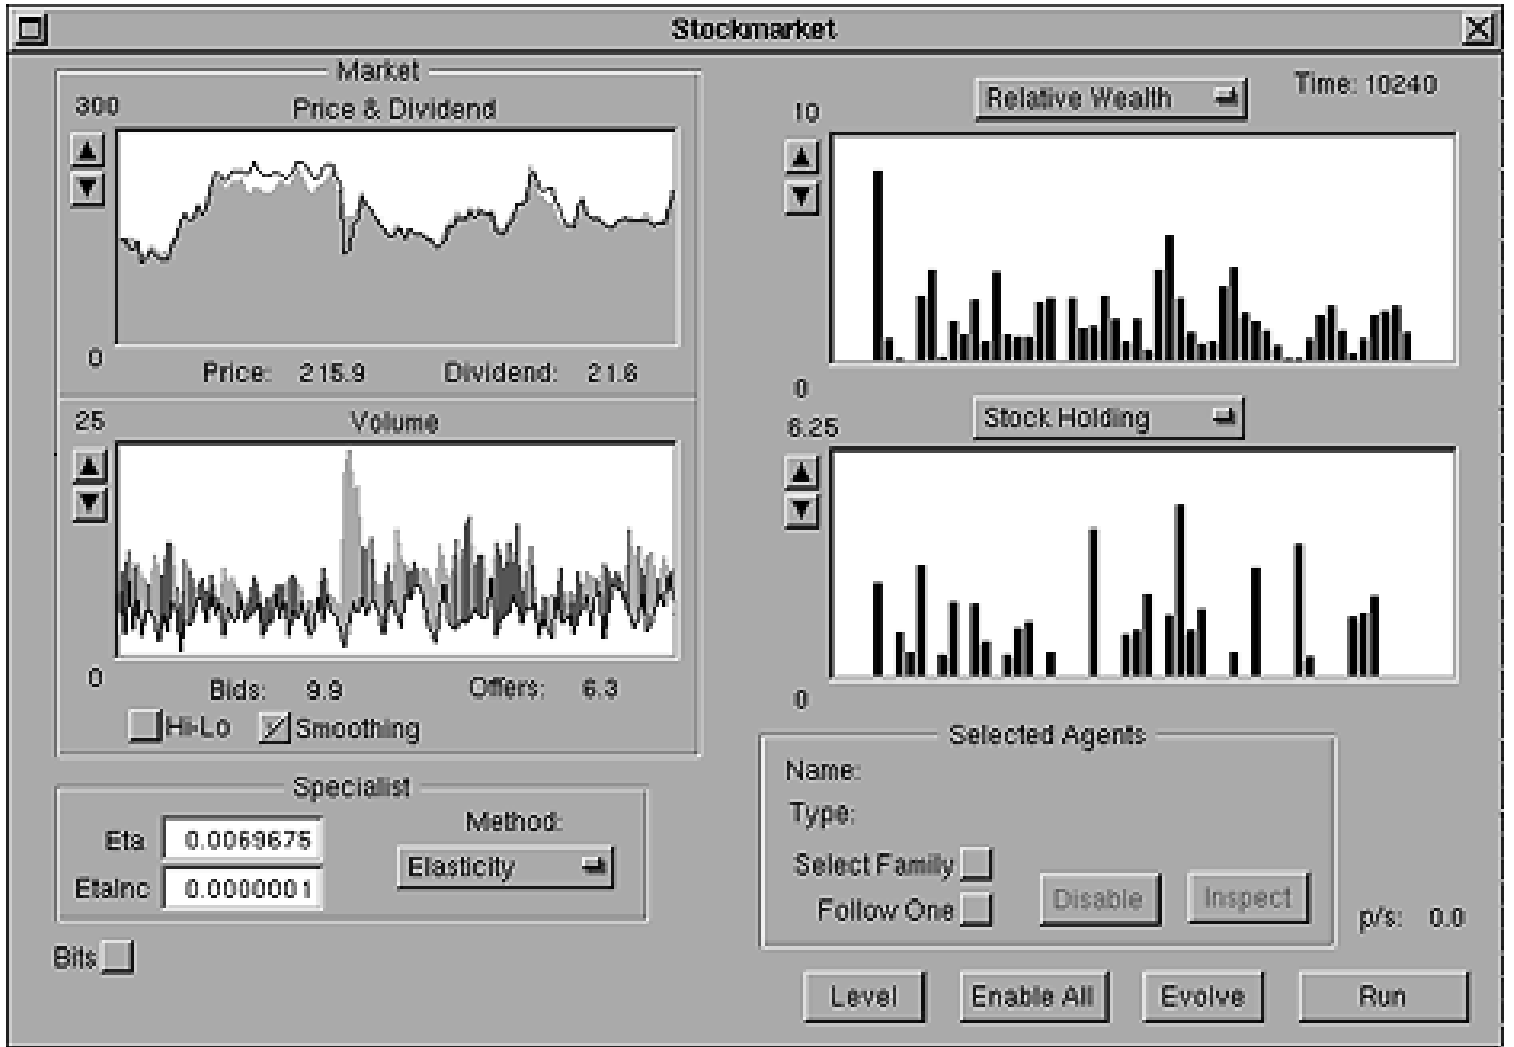


Figure 10. A frozen moment in the surrogate stock market.

#### IV. THE ROLE OF CHAOS AND FRACTALS

One of the most pernicious misconceptions about complex systems is that complexity and chaotic behavior are synonymous. On the basis of the foregoing discussion of emergence, it's possible to put the role of chaos in complex systems into proper perspective. Basically, the situation is that if one focuses attention on the time evolution of an emergent property, such as the price movements of a stock or the daily changes in temperature, then that property may well display behavior that is completely deterministic, yet is indistinguishable

from a random process; in other words, it is “chaotic”. So chaos is an epiphenomenon, one that often goes hand-in-hand with complexity—but does not necessarily follow from it in all cases. What is important from a system-theoretic standpoint is what is happening at the lower level of the agents—traders, drivers, molecules—whose interaction create the emergent patterns, not the possibly chaotic behavior that these patterns display. Here is an example illustrating the difference.

El Farol is a bar in Santa Fe, at which Irish music used to be played every Thursday evening. As a result, the same Irish economist, W. Brian Arthur, who created the artificial stock market described above, was fond of going to the bar to hear this group each Thursday. But he wasn’t fond of doing so in the midst of a madhouse of pushing-and-shoving drinkers. So Arthur’s problem each Thursday was to decide whether the crowd at El Farol would be so large that the enjoyment he received from the music would be outweighed by the irritation of having to listen to the performance drowned in the shouts, laughter, and raucous conversation. Arthur attacked the question of whether or not to attend in analytical terms by constructing a simple model of the situation.

Assume, said Arthur, that there are 100 people in Santa Fe, each of whom would like to listen to the music. But none of them wants to go if the bar is going to be too crowded. To be specific, suppose that all 100 people know the attendance at the bar for each the past several weeks. For example, such a record might be ... 44, 78, 56, 15, 23, 67, 84, 34, 45, 76, 40, 56, 23, and 35 attendees. Each individual then independently employs some prediction procedure to estimate how many people will appear at the bar on the coming Thursday evening. Typical predictors of this sort might be:

- a. the same number as last week (35);
- b. a mirror image around 50 of last week's attendance (65);
- c. a rounded-up average of the past four weeks' attendance (39);
- d. the same as two weeks ago (23).

Suppose each person decides independently to go to the bar if his or her prediction is that fewer than 60 people will go; otherwise, the person stays home. In order to make this prediction, every person has his or her own individual set of predictors and uses the currently most accurate one to forecast the coming week's attendance at El Farol. Once each person's forecast and decision to attend has been made, people converge on the bar, and the new attendance figure is published the next day in the newspaper. At this time, everyone updates the accuracies of all of the predictors in his or her particular set, and things continue for another round. This process creates what might be termed an "ecology" of predictors.

The problem faced by each person is then to forecast the attendance as accurately as possible, knowing that the actual attendance will be determined by the forecasts others make. This immediately leads to an "I-think-you-think-they-think- . . . " -type of regress—a regress of a particularly nasty sort. For suppose that someone becomes convinced that 87 people will attend. If this person assumes others are equally smart, then it's natural to assume they will also see that 87 is a good forecast. But then they all stay home, negating the accuracy of that forecast! So no shared, or common, forecast can possibly be a good one; in short, deductive logic fails. So, from a scientific point of view, the problem comes

down to how to create a *theory* for how people decide whether or not to turn up at El Farol on Thursday evening and for the dynamics that these decisions induce.

It didn't take Arthur long to discover that it seems to be very difficult even to formulate a useful model of this decision process in conventional mathematical terms. So he decided to turn his computer loose on the problem and to create the "would-be" world of El Farol inside it in order to study how electronic people would act in this situation. What he wanted to do was look at how humans reason when the tools of deductive logic seem powerless to offer guidelines as to how to behave. As an economist, his interest is in self-referential problems—situations in which the forecasts made by economic agents act to *create* the world they are trying to forecast. Traditionally, economists look at such worlds using the idea of *rational expectations*. This view assumes homogeneous agents, who agree on the same forecasting model and know that others know that others know that . . . they are using this forecast model. The classical view then asks which forecasting model would be consistent, on the average, with the outcome that it creates. But nobody asks how agents come up with this magical model. Moreover, if you let the agents differ in the models they use, you quickly run into a morass of technical—and conceptual—difficulties.

What Arthur's experiments showed is that if the predictors are not too simplistic, then the number of people who will attend fluctuates around an average level of 60. And, in fact, whatever threshold level Arthur chose, that level always seemed to be the long-run average of the number of attendees. In addition, the computational experiments turned up an even more intriguing pattern—at least for mathematical system theorists: The number of people going to the bar each week is a purely deterministic function of the individual

predictions, which themselves are deterministic functions of the past number of attendees. This means that there is no inherently random factor dictating how many people actually turn up. Yet the outcome of the computational experiments suggests that the actual number going to hear the music in any week looks more like a random process than a deterministic function. The graph in Figure 11 shows a typical record of attendees for a 100-week period when the threshold level is 60.




Figure 11. Attendance at the bar over a 100-week period.

These experimental observations lead to a fairly definite and specific mathematical conjecture:

- Under suitable conditions (to be determined) on the set of predictors, the average number of people who actually go to the bar converges to the threshold value as the number of time periods becomes large.

System theorists will also want to ponder the related conjecture having to do with chaos:

- Under the same set of suitable conditions on the sets of predictors, the time-series of attendance levels is a deterministically-random process, that is, it's "chaotic."

So here we see a prototypical situation of a complex, adaptive system in which a global property (number of attendees at the bar) emerges from the individual decisions of lower-level agents (the 100 Irishmen). Moreover, the temporal fluctuation of the emergent property appears to display chaotic behavior. But it's not hard to change the situation so as to remove the chaos. For example, if every Irishman always uses the rule, "Go to the bar regardless of the past record of attendees", then the record of attendees stays fixed at 100 for all time. This is certainly *not* chaotic behavior.

Fractals are geometric objects closely associated with chaotic processes. The relationship is rather straightforward. One of the most evident features of a chaotic dynamical system is the property of sensitivity to initial positions. This means that starting the same process from two different initial states as close together as one might please, leads to long-run behavior that may be totally different. The famous "butterfly effect", in which a butterfly flapping its wings in China today can give rise to tornados in Kansas tomorrow is a good illustration of this phenomenon.

The pathological sensitivity of chaotic systems to small disturbances in their initial state means that the space of initial states can be divided up into regions, such that two states in the same region lead have the same long-run behavior. This gives a geometrical decomposition of the space of initial states. If the system is not chaotic, this geometry is usually rather straightforward, with the borders between different regions being described by simple, smooth curves like arcs of circles or even straight lines. But when the dynamical

system is chaotic, the curves separating one region from another are highly irregular objects termed “fractal”. The well-chronicled Mandelbrot set is certainly the most famous of these objects. But recent work in engineering, physics, biology, and other areas has shown the ubiquity of fractals throughout nature.

What we can conclude from this discussion is that chaotic behavior and fractal geometries are properties of patterns that emerge from lower-level activities of agents making up a complex, adaptive system. So just as one should not confuse the map with the territory, one should also not confuse chaos and fractals with complexity. They are just not the same thing at all.

## V. THE SCIENCE OF COMPLEXITY

Recall in the discussion above of the El Farol Problem that the Irishmen faced the question of using their currently best rule to estimate how many music lovers would appear at the bar in the coming week. On the basis of this prediction, each individual then chose to go to the bar or stay home, with the total number attending being reported to everyone in the following week. At that time, each Irishman revised his or her set of predictors, using the most accurate predictor to estimate the attendance in the coming week. The key components making up the El Farol Problem are exactly the key components in each and every complex, adaptive system, and a decent mathematical formalism to describe and analyze the El Farol Problem would go a long way toward the creation of a viable theory of such processes. These key components are:

*A medium-sized number of agents*—The El Farol Problem postulates 100 Irishmen, each of whom acts independently in deciding to go or not go to the bar on Thursday



evening. In contrast to simple systems—like superpower conflicts, which tend to involve a small number of interacting agents—or large systems—like galaxies or containers of gas, which have a large enough collection of agents that we can use statistical means to study them—complex systems involve what we might call a medium-sized number of agents. Just like Goldilocks’s porridge, which was not too hot and not too cold, complex systems have a number of agents that is not too small and not too big, but just right to create interesting patterns of behavior.

*Intelligent and adaptive agents*—Not only are there a medium-sized number of agents, these agents are intelligent and adaptive. This means that they make decisions on the basis of rules, and that they are ready to modify the rules they use on the basis of new information that becomes available. Moreover, the agents are able to generate new rules that have never before been used, rather than being hemmed in by having to choose from a set of preselected rules for action. This means that an ecology of rules emerges, one that continues to evolve during the course of the process.

*Local information*—In the real world of complex systems, no agent knows what *all* the other agents are doing. At most, each person gets information from a relatively small subset of the set of all agents, and processes this “local” information to come to a decision as to how he or she will act. In the El Farol Problem, for instance, the local information is as local as it can be, because each Irishman knows only what he or she is doing; none have information about the actions taken by any other agent in the system. This is an extreme case, however, and in most would-be worlds the agents are more like drivers in a

transport system or traders in a market, each of whom has information about what a few of the other drivers or traders are doing.

So these are the components of all complex, adaptive systems like the El Farol situation—a medium-sized number of intelligent, adaptive agents interacting on the basis of local information. At present, there appears to be no known mathematical structures within which we can comfortably accommodate a description of the El Farol Problem. This suggests a situation completely analogous to that faced by gamblers in the seventeenth century, who sought a rational way to divide the stakes in a game of dice when the game had to be terminated prematurely (probably by the appearance of the police or, perhaps, the gamblers' wives). The description and analysis of that very definite real-world problem led Fermat and Pascal to the creation of a mathematical formalism we now call probability theory. At present, complex-system theory still awaits its Pascal and Fermat. The mathematical concepts and methods currently available were developed, by and large, to describe systems composed of material objects like planets and atoms. It is the development of a proper theory of complex systems that will be the capstone of the transition from the material to the informational.

## References

1. Casti, J. *Reality Rules: Picturing the World in Mathematics I—The Fundamentals, II—The Frontier*. New York: Wiley, 1992 (paperback edition, 1997).
2. Casti, J. *Complexification*. New York: HarperCollins, 1994.
3. Gleick, J. *Chaos*. New York: Viking, 1987.

4. Stewart, I. *Does God Play Dice?* Oxford: Basil Blackwell, 1989.
5. Jackson, E. Atlee. *Perspectives of Nonlinear Dynamics*. Vols. 1&2. Cambridge: Cambridge University Press, 1990.
6. Schroeder, M. *Fractals, Chaos, Power Laws*. New York: W. H. Freeman, 1991.
7. Mandelbrot, B. *The Fractal Geometry of Nature*. San Francisco: W. H. Freeman, 1982.
8. Peitgen, H.-O., D. H. Jürgens, and D. Saupe. *Fractals for the Classroom*. Parts 1 & 2. New York: Springer, 1992.
9. Nicolis, J. *Chaos and Information Processing*. Singapore: World Scientific, 1991.
10. Casti, J. *Would-Be Worlds*. New York: John Wiley & Sons, 1997.