

The Simply Complex

John L. Casti

THE FINGERPRINTS OF COMPLEXITY

In everyday parlance the term “complex” is generally used to describe a person or thing that is composed of many interacting components whose behavior and/or structure is just plain hard to understand. The behavior of national economies, the human brain and a rain forest ecosystem are all good examples of complex systems. These examples underscore the point that sometimes a system may be structurally complex, like a mechanical clock, but behavior 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 both in structure and behavior.

As these examples indicate, 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 more importantly the tools, to study such systems in a controlled, repeatable, scientific fashion. And hopefully 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 complex systems like a national economy or the human




Figure 1. A toy rotator with unpredictable behavior.

immune system because it was simply impractical, too expensive—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 the arrival of today’s computers, we can actually build complete silicon surrogates of these systems inside our computing machines, and use these “would-be worlds” as laboratories within which to look at the workings—and behaviors—of the complex systems of everyday life.

To speak of a system as being *complex* suggests identifying features separating complex systems from those that are in some sense simple. Here are a few of the most important of these “fingerprints” of complexity.

- *Instability:* Complex systems tend to have many possible modes of behavior, often shifting between these modes as the result of small changes in some factors governing the system. For instance, the flow of water or oil through a pipe is smooth when the flow velocity is low. But if the velocity is increased beyond a critical level (that depends on the

viscosity of the fluid), eddies and whirlpools appear. And if the velocity is increased still further, the frothy, chaotic motion of fully-developed turbulence sets in.

- *Irreducibility:* Complex systems come as a unified whole; they cannot be studied by breaking them into their component parts and looking at the parts in isolation. The behavior of the system is determined by the interaction among the parts, and any “tearing” of the system into pieces destroys the very aspects that give it its individual character. A good illustration of this is the well-known problem of protein folding. Every protein is formed as a chain of amino acids, strung together like beads on a necklace. Once the protein has been assembled, it folds up into a unique three-dimensional configuration that determines its function in the living organism, a configuration that is determined completely by the one-dimensional sequence of amino acids. But in order to know what this final configuration will be, it is simply not possible to separate the question of how the protein will fold into a set of smaller subproblems by cutting the protein at various spots and seeing how these subchains of amino acids fold, and then cementing together somehow the solutions of these individual subproblems. To see how the protein will fold it must be studied as a single, integrated whole.

- *Adaptability:* Complex systems tend to be composed of many intelligent agents, who take decisions and actions on the basis of partial information about the entire system. Moreover, these agents are capable of changing their decision rules on the basis of such information. A driver in a road-traffic network or a trader in a financial market illustrate this point nicely, since in both cases the agent receives partial information about the system he or she is a part of—traffic conditions for the driver, prices and market trends for the

trader—and takes actions on the basis of this information. As a result of these actions, both the driver and the trader gain information about what the rest of the system—the other drivers or traders—are doing. The agent can then modify his or her decision rules accordingly. In short, complex systems generally have the capability of learning about their environment and changing their behavioral responses to it in the light of new information.

In passing, let me note that human beings are not the only type of agent that fits this mold. Molecules, corporations and living cells also qualify as intelligent, adaptive agents, who can change their behavior in response to changes in their environment.

- *Emergence:* 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. Everyday tap water illustrates the general idea, as its component parts—hydrogen and oxygen—are both highly flammable gases yet combine to produce a compound that is neither. Thus, the properties of being a liquid and noncombustible are emergent properties arising from the interaction of the hydrogen and oxygen “agents.” A similar 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.

With these ideas in mind, let's turn to some examples of systems displaying these sorts of features.

AN ARTIFICIAL STOCK MARKET

Around 1988, Brian Arthur, an economist from Stanford, and John Holland, a computer scientist from the University of Michigan, were sharing a house in Santa Fe while both were visiting the Santa Fe Institute. During endless hours of evening conversations over many bottles of wine, Arthur and Holland 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 not-very-surprising 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 recruited physicist Richard Palmer, finance theorist Blake LeBaron and market trader Paul Tayler to help them construct their electronic market, where they could, in effect, play god by manipulating trader's strategies, market parameters and all the other things that cannot be done with real 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 below in Figure 2.

A typical moment in this artificial market is displayed in Figure 3. 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.

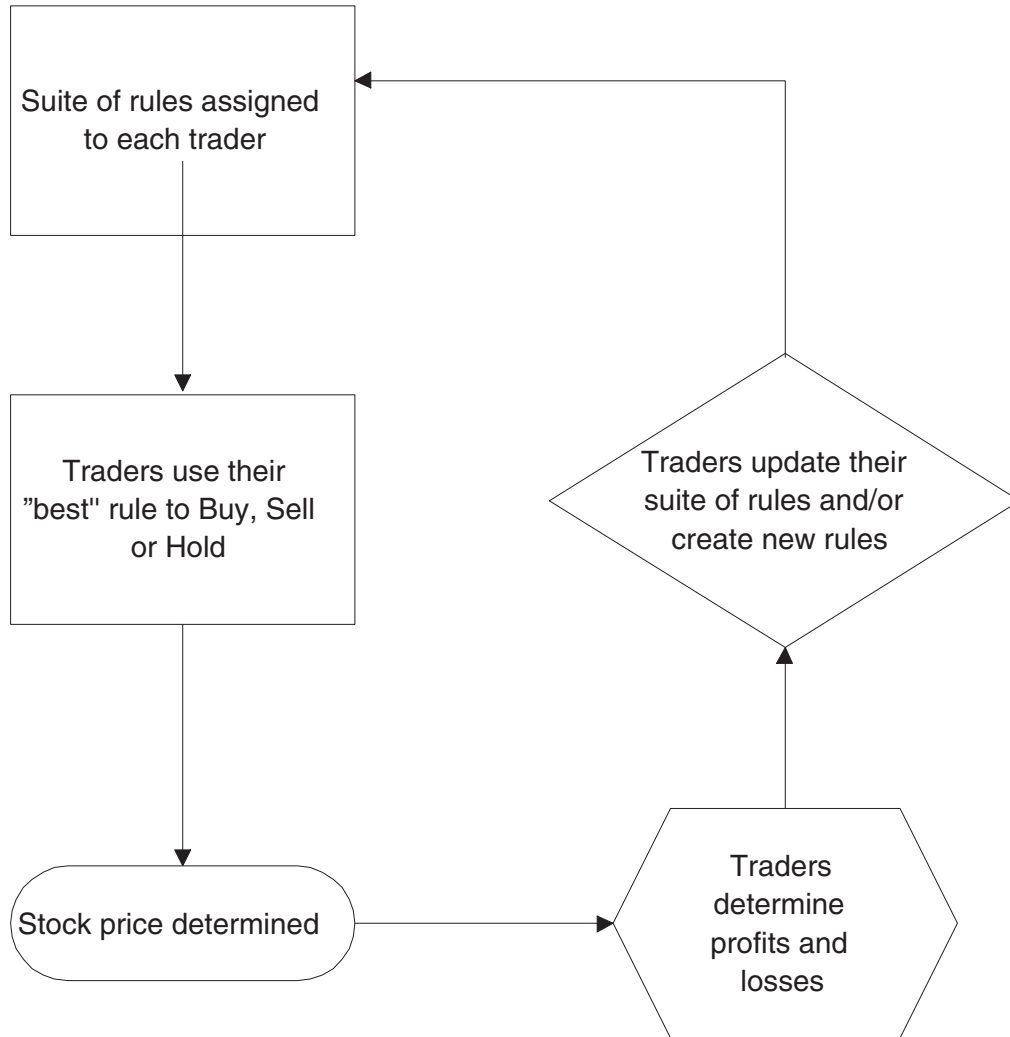


Figure 2. The logical flow of activity in the stock market.

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.

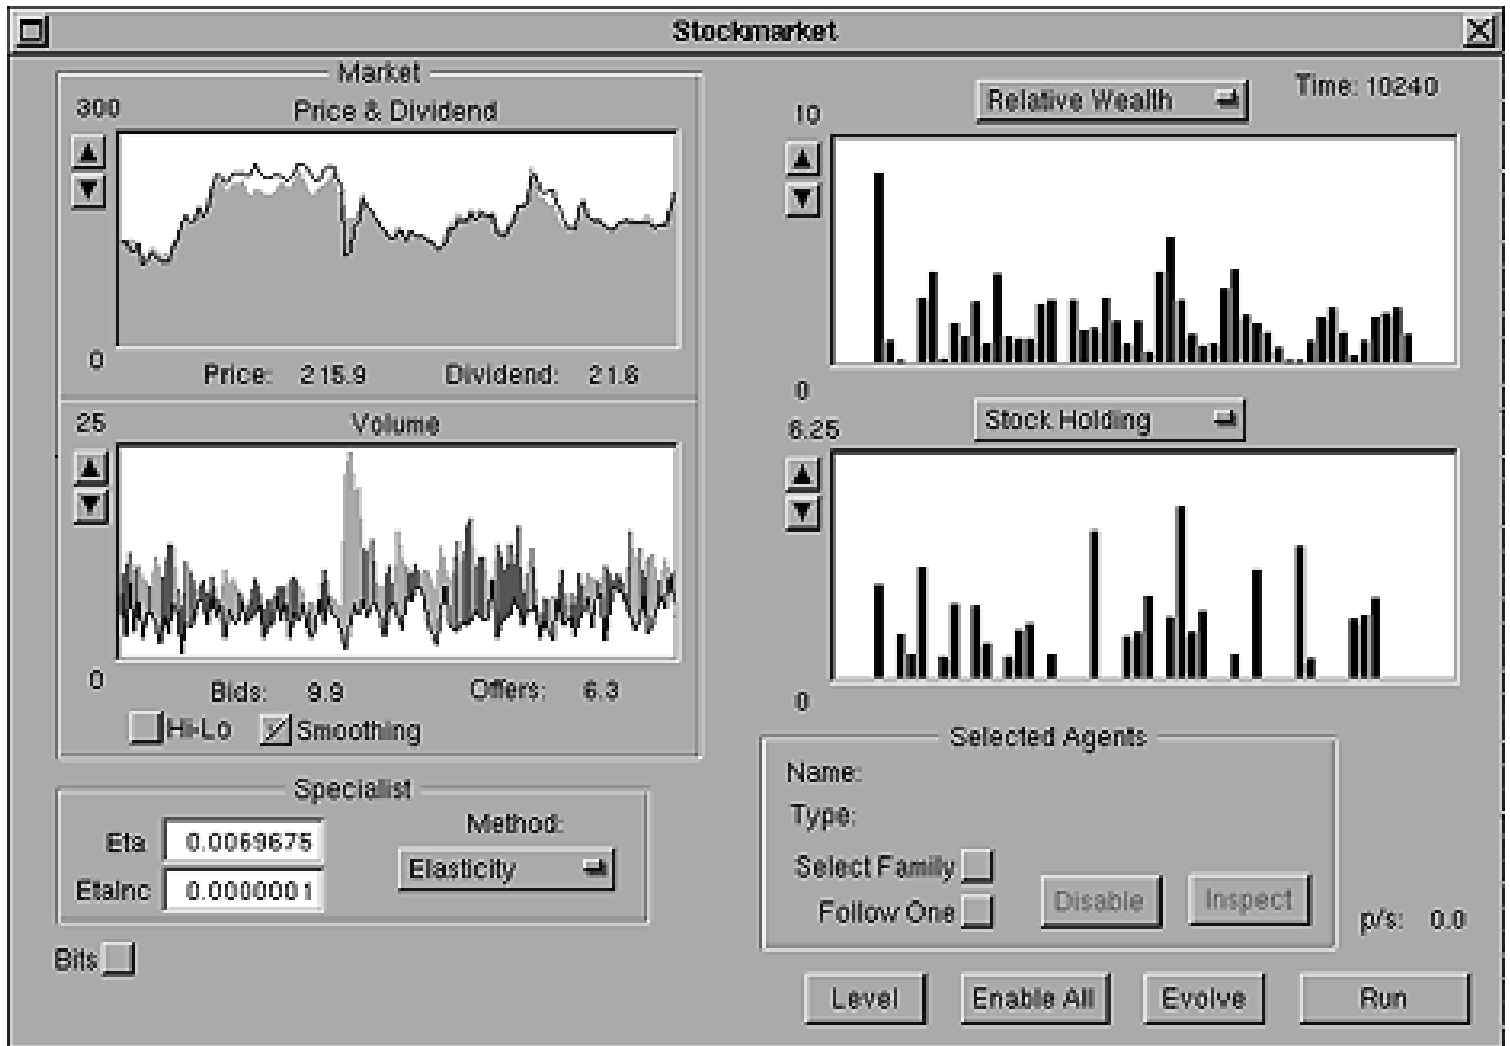


Figure 3. A frozen moment in the surrogate stock market.

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.

Let's now turn to the realm of biology and look at another example of complex adaptive systems in action.

THE ARRIVAL OF THE FITTEST

Certainly among the most refractory questions that science addresses are the so-called “origins” problems, involving issues like the origin of the universe, origin of life, origin of language and the origin of humans. The difficulty, of course, is that these are one-time-only events, while the ability to perform controlled, repeatable experiments is part and parcel of the scientific method. But the computer has changed our way of looking at such questions, allowing just such types of experiments to be performed on origin questions in surrogate worlds rather than for the real thing.

Recently, Walter Fontana of the University of Vienna and Leo Buss of Yale University asked themselves the following question: Suppose we could create a duplicate of the Earth of 4 billion years ago and set it running. What kinds of biological organizations and structures would we expect to see today in that “second Earth”? Put more prosaically, what would be conserved if the tape were played twice? Could we reasonably expect to see things like zebras, ant colonies and platypuses emerge again? Or would the vagaries of evolution cause this second Earth show species and biological organizations totally different from what we see on Earth today? In short, what is contingent and what is necessary about life?

To address this question, Fontana and Buss found inspiration in the structures of theoretical computer science. In particular, they realized that biological entities be they the molecules of life or full-scale organisms are not simply passive objects shoved around by the laws of natural selection. Rather, every such entity is simultaneously such a passive object *and* an active operator capable of acting on other such objects to create/destroy

other such objects. Computer science has dealt with such entities for years using a formal structure called the lambda calculus, a logical framework providing rules for how such new entities are created and for how to simplify them to their so-called “normal forms.”

In their experiments, Fontana and Buss created a primordial soup of such entities, letting several hundred of them collide and interact randomly within a kind of electronic test tubs. The rules of the lambda calculus were employed to determine the end results of various random interactions, as well as to single out which new entities would survive to succeeding generations. And what were the results?

At first, something rather disappointing happens. After many collisions, the entities in the soup are all the same, consisting of particles that simply replicate themselves. Fontana and Buss discovered that this “genetic takeover” is due simply to the possibility for objects to build copies of themselves directly, without needing the help of any other entities in the system. So to get something interesting to come crawling up out of the soup, they introduced a prohibition against this type of autonomous self-reproduction. Analysis after many thousands of collisions in the soup now yields quite a different outcome.

The soup contains many different types (species?) of entities, a substantial fraction of which is involved in a self-sustaining network of mutual production pathways. Moreover, all the entities in the system share a common structure that conforms to specific patterns representable as a kind of “biological grammar.” Thus, the system is very far from being the original random mixture of entities. And, finally, the transformation relationships among all the entities obey a small number of laws. Taken together, these properties constitute an answer to the question of what would be conserved if the tape were played twice. What we would expect to see is not specific objects but rather specific types of

organizational structures of entities. So while zebras are unlikely to appear in this second Earth, the kind of family and ecosystem structures of which zebras are a part would very likely reemerge—but with some other kind of animal playing the role that zebras play in today’s Earth.

WHITHER HENCE, COMPLEXITY?

Comparing the artificial stock market and the electronic ecosystems, we see the power of complexity theory to address *experimentally* questions that heretofore science has been powerless to explore, questions like how various “fashions” come and go in a financial market or what kind of organizational structures we might expect to see on a “second Earth” in the Andromeda galaxy. The basic principle at work here is that the history of physics tells us that the ability to do laboratory experiments is a necessary step in the creation of scientific theory of any sort of phenomenon. Complex system theorists are now in the situation that physicists were in at the time of Galileo. For the first time in history we can actually perform experiments on genuine complex systems. And with this capability, can a decent theory of such structures be far behind?

References

- Casti, J. *Complexification*. HarperCollins, New York, 1994.
- Casti, J. *Would-Be Worlds*. Wiley, New York, to appear fall 1996.
- Rieck, C. Evolutionary Simulation of Asset Trading Strategies. *Many-Agent Simulation and Artificial Life*, E. Hillebrand and J. Stender, eds., IOS Press, Amsterdam, 1994, pp. 112–136.

Arthur, W. B. et al. An Artificial Stock Market. Santa Fe Institute Working Paper, Santa Fe, NM, in press.

Fontana, W. and L. Buss. What Would be Conserved if “the Tape Were Played Twice”? *Proc. Nat. Acad. Sci. USA*, 91: 757–761, 1994.