Complexity and the Economy

W. Brian Arthur

Citibank Professor, Santa Fe Institute

Paper appeared in *Science*, 2 April 1999, **284**, 107-109

Common to all studies on complexity are systems with multiple elements adapting or reacting to the pattern these elements create. The elements might be cells in a cellular automaton, or ions in a spin glass, or cells in an immune system, and they may react to neighboring cells' states, or local magnetic moments, or concentrations of B and T cells—"elements" and the "patterns" they respond to vary from one context to another. But the elements adapt to the world—the aggregate pattern—they co-create. Time enters naturally here via adjustment and change: as the elements react, the aggregate changes, as the aggregate changes, elements react anew. Barring some asymptotic state or equilibrium reached, complex systems are systems in process, systems that constantly evolve and unfold over time.

Such systems arise naturally in the economy. Economic agents, be they banks, consumers, firms, or investors, continually adjust their market moves, buying decisions, prices, and forecasts to the situation these moves or decisions or prices or forecasts together create. But unlike ions in a spin glass which always react in a simple way to their local magnetic field, economic "elements"—human agents—react with strategy and foresight by considering outcomes that might result as a consequence of behavior they might undertake. This adds a layer of complication to economics not experienced in the natural sciences.

Conventional economic theory chooses not to study the unfolding of the patterns its agents create, but rather to simplify its questions in order to seek analytical solutions. Thus it asks what behavioral elements (actions, strategies, expectations) are consistent with the aggregate patterns these behavioral elements co-create? For example, general equilibrium theory asks: what prices and quantities of goods produced and consumed are consistent with—would pose no incentives for change to—the overall pattern of prices and quantities in the economy's markets. Game theory asks: what strategies, moves, or allocations are consistent with—would induce no further reactions to—the potential outcomes these strategies, moves, allocations might imply. Rational expectations economics asks: what forecasts (or expectations) are consistent with—are on average validated by—the outcomes these forecasts and expectations together create. Conventional economics thus studies consistent patterns—patterns in behavioral equilibrium, patterns that would induce no further reaction. Economists at the Santa Fe Institute, Stanford, MIT, Chicago, and other institutions, are now broadening this equilibrium approach by turning to the question of how actions, strategies, or expectations might react in general to—might endogenously change with—the aggregate patterns these create [1]. The result, complexity economics, is not an adjunct to standard economic theory, but theory at a more general, out-of-equilibrium level.

The type of systems I have described become especially interesting if they contain nonlinearities in the form of positive feedbacks. In economics positive feedbacks arise from increasing returns [2] [3]. To ensure a unique, predictable equilibrium is reached, standard economics usually assumes diminishing returns. If one firm gets too far ahead in the market, it runs into higher costs or some other negative feedback and the market is shared at a predictable, unique equilibrium. When we allow positive feedbacks, or increasing returns, a different outcome arises. Consider the market for online services of a few years back, in which three major companies competed: Prodigy, Compuserve, and America Online. As each gained in membership base it could offer a wider menu of services as well as more members to share specialized hobby and chatroom interests with—there were increasing returns to expanding the membership base. Prodigy was first in the market, but by chance and strategy American Online got far enough ahead to gain an unassailable advantage. Today it dominates. Under different circumstances, another rival might have taken the market. Notice the properties here: a multiplicity of potential "solutions"; the outcome actually reached is not predictable in advance; it tends to be locked in; it is not necessarily the most efficient economically; it is subject to the historical path taken; while the companies may start equal, the outcome is asymmetrical. These properties have counterparts in non-linear physics where similar positive feedbacks are present. What economists call multiple equilibria, non-predictability, lock-in, inefficiency, historical path dependence, and asymmetry; physicists call multiple meta-stable states, unpredictability, phase- or mode-locking, high-energy ground states, non-ergodicity, and symmetry breaking [3].

Increasing returns problems have been discussed in economics for a long time. A hundred years ago, Alfred Marshall [5] noted that if firms gain advantage as their market share increases, "whatever firm first gets a good start will obtain a monopoly." But the conventional, static equilibrium approach gets stymied by indeterminacy: If there is a multiplicity of equilibria, how might one be reached? The process-oriented, complexity approach suggests a way to deal with this. In the actual economy, "small random events" happen—in the online-services case "random" interface improvements, new offerings, word-of-mouth recommendations. Over time increasing returns magnifies the cumulation of such events to "select" the outcome randomly. Thus increasing returns problems in economics are best seen as dynamic processes with random events and natural positive feedbacks—as nonlinear stochastic processes. This shift from a static outlook into a process orientation is common to complexity studies. Increasing returns problems are being studied intensively in market allocation theory [3], international trade theory [6], the evolution of technology choice [7], economic geography [8], and the evolution of patterns of poverty and segregation [9]. The common finding that economic structures can crystallize around small events and lock in is beginning to change policy in all these areas toward an awareness that governments should avoid both extremes of coercing a desired outcome or keeping strict hands off, and instead seek to push the system gently toward favored structures that can grow and emerge naturally. Not a heavy hand, not an invisible hand, but a nudging hand.

Once we adopt the complexity outlook, with its emphasis on the formation of structures rather than their given existence, problems involving prediction in the economy look different. The conventional approach asks what forecasting model (or expectations) in a particular problem, if given and shared by all agents, would be consistent with—would be on average validated by—the actual time series this forecasting model would in part generate. This "rational expectations" approach is valid. But it assumes that agents can somehow deduce in advance what model will work, and that everyone "knows" that everyone knows to use this model (the *common knowledge* assumption.) What happens when forecasting models are not obvious and must be formed individually by agents who are not privy to the expectations of others?

Consider as an example my El Farol Bar Problem [10]. One hundred people must decide independently each week whether to show up at their favorite bar (*El Farol* in Santa Fe). The rule is that if a person predicts that more that 60 (say) will attend, he will avoid the crowds and stay home; if he predicts fewer than 60 he will go. Of interest are how the bar-goers each week might predict the numbers showing up, and the resulting dynamics of the numbers attending. Notice two features of this problem. Our agents will quickly realize that predictions of how many will attend depend on others' predictions of how many attend (because that determines their attendance). But others' predictions in turn depend on their predictions of others' predictions. Deductively there is an infinite regress. No "correct" expectational model can be assumed to be common knowledge, and from the agents' viewpoint, the problem is ill-defined. (This is true for most expectational problems, not just for this example.) Second, and diabolically, any commonalty of expectations gets broken up: If all use an expectational model that predicts few will go, all will go, invalidating that model. Similarly, if all believe most will go, nobody will go, invalidating that belief. Expectations will be forced to differ.

In 1993 I modeled this situation by assuming that as the agents visit the bar, they act inductively—they act as statisticians, each starting with a variety of subjectively chosen expectational models or forecasting hypotheses. Each week they act on their currently most accurate model (call this their *active* predictor). Thus agents' beliefs or hypotheses compete for use in an *ecology* these beliefs create. Computer simulation (Fig. 1) showed that the mean attendance quickly converges to 60. In fact, the predictors self-organize into an equilibrium "ecology" in which of the active predictors 40% on average are forecasting above 60, 60% below 60. This emergent ecology is organic in nature. For, while the population of active predictors splits into this 60/40 average ratio, it keeps changing in membership forever. Why do the predictors self-organize so that 60 emerges as average attendance and forecasts split into a 60/40 ratio? Well, suppose 70% of predictors forecasted above 60 for a longish time, then on average only 30 people would show up. But this would validate predictors that forecasted close to 30, restoring the "ecological" balance among predictions. The 40%–60% "natural" combination becomes an emergent structure. The Bar Problem is a miniature expectational economy, with complex dynamics. [11].

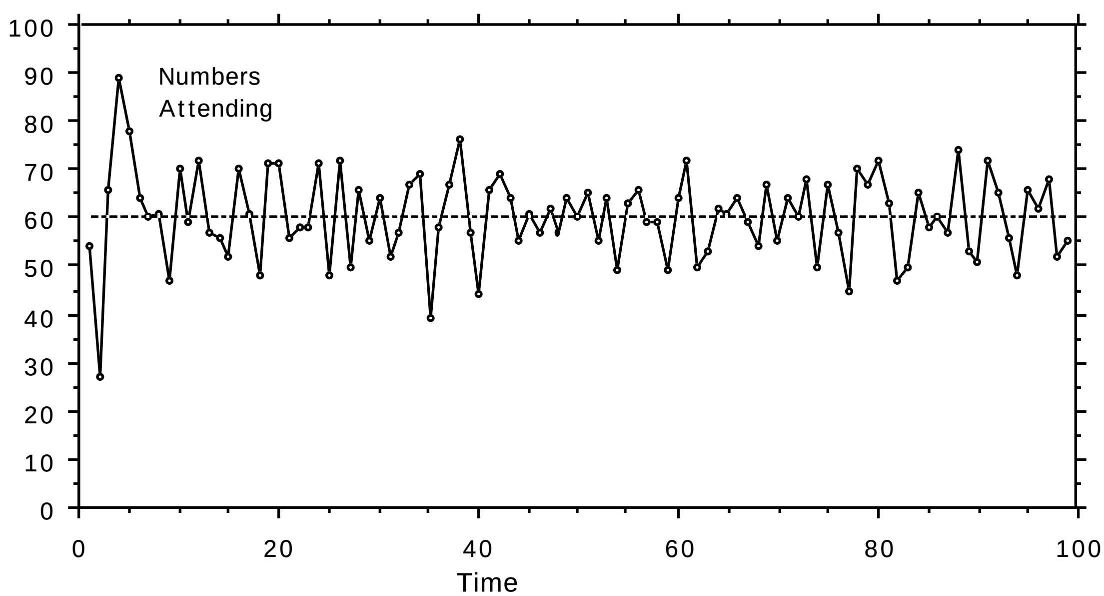

Figure 1. Bar Attendance in the first 100 Weeks.

One important application of these ideas is in financial markets. Standard theories of financial markets assume rational expectations—that agents' adopt uniform forecasting models that are on average validated by the prices these forecast [12]. The theory works well to first order. But it doesn't account for actual market "anomalies" such as unexpected price bubbles and crashes, random periods of high and low volatility (price variation), and the heavy use of technical trading (trades based on the recent history of price patterns). Holland, LeBaron, Palmer, and I [13] have created a model which relaxes rational expectations by assuming, as in the Bar Problem, that investors cannot assume or deduce expectations but must discover them. Our agents continually create and use multiple "market hypotheses"—individual, subjective, expectational models—of future prices and dividends within an artificial stock market on the computer. These "investors" are individual, artificially-intelligent computer programs that can generate and discard expectational "hypotheses," and make bids or offers based on their currently most accurate hypothesis. The stock price forms from their bids and offers, and thus ultimately from agents' expectations. So this market-in-the-machine is its own self-contained, artificial financial world. Like the bar, it is a "mini-ecology" in which expectations compete in a world these expectations create.

Within this computerized market, we found two phases or regimes. If parameters are set so that our artificial agents update their hypotheses slowly, the diversity of expectations collapses quickly into homogeneous rational expectations ones. The reason is that if a majority of investors believes something close to the rational expectations forecast, then resulting prices will validate it, and deviant or mutant predictions that arise in the population of expectational models will be rendered inaccurate. Standard finance theory, under these special circumstances, is upheld. But if the rate of updating of hypotheses is turned up, the market undergoes a phase transition into a "complex regime" and displays several of the "anomalies" observed in real markets. It develops a rich "psychology" of divergent beliefs that don't converge over time. Expectational rules such as "If the market is trending up, predict a 1% price rise" that appear randomly in the population of hypotheses can become mutually reinforcing—if enough investors act on these, the price will indeed go up. Thus sub-populations of mutually reinforcing expectations arise, agents bet on these (therefore technical trading emerges) and this causes occasional bubbles and crashes. Our artificial market also shows periods of high volatility in prices followed randomly by periods of low volatility. This is because if some investors "discover" new, profitable hypotheses, they change the market slightly, causing other investors to also change their expectations. Changes in beliefs therefore ripple through the market in avalanches of all sizes, causing periods of high and low volatility. We conjecture that actual financial markets, which show exactly these phenomena, lie in this "complex" regime.

Conclusion

After two centuries of studying equilibria—static patterns that call for no further behavioral adjustments—economists are beginning to study the general emergence of structures and unfolding of patterns in the economy. Complexity economics is not a temporary adjunct to static economic theory, but theory at a more general, out-of-equilibrium level. The approach is making itself felt in every area of economics: game theory [14], the theory of money and finance [15], learning in the economy [16], economic history [17], the evolution of trading networks [18], the stability of the economy [19], and political economy [20]. It is helping us understand phenomena such as market instability, the emergence of monopolies, and the persistence of poverty in ways that will help us deal with these. And it is bringing an awareness that policies succeed better by influencing the natural processes of formation of economic structures, than by forcing static outcomes.

When viewed in out-of-equilibrium formation, economic patterns sometimes simplify into the simple, homogeneous equilibria of standard economics. More often they are ever-changing, showing perpetually novel behavior and emergent phenomena. Complexity therefore portrays the economy not as deterministic, predictable and mechanistic; but as process-dependent, organic and always evolving.

References and Notes

1. See the general collections in P. Anderson, K. J. Arrow, D. Pines, Eds. *The Economy as an Evolving Complex System* (Addison-Wesley, Reading, Mass., 1988) and W. B. Arthur, S. N. Durlauf, D. A. Lane, Eds. *The Economy as an Evolving Complex System II* (Addison-Wesley, Reading, Mass., 1997).

2. W. B. Arthur, *Scientific American*, 92, (1990).

3. W. B. Arthur, *Increasing Returns and Path Dependence in the Economy* (University of Michigan Press, Ann Arbor, 1994).

4. Note that I have avoided exact definitions of "complexity" and "complex systems." Technically, the systems I have described are referred to as *adaptive nonlinear networks* (J. H. Holland's term), and typically if they exhibit certain properties that have to do with the multiplicity of potential patterns or with the coherence or propagation of sub-structures they are said to be "complex." Definitions vary widely.

5. A. Marshall, *Principles of Economics*, 459, (8th Edition, Macmillan, London, 1920).

6. E. Helpman and P. R. Krugman, *Market Structure and Foreign Trade*, (MIT Press, Cambridge, Mass., 1985).

7. W. B. Arthur, *Econ. Journal*, 99, 116, (1989).

8. W. B. Arthur in *Math. Social Sciences*, 19, 235, (1990); P. R. Krugman in *J. Pol. Econ.* 99, 483, (1991) and in *The Economy as an Evolving Complex System II*, (Op. Cit.); P. R. Krugman, *Geography and Trade* (MIT Press, Cambridge, Mass, 1991).

9. S. N. Durlauf, in The Economy as an Evolving Complex System II, (Op. Cit.) and in J. Econ. Growth 1, 75, (1996).

10. W. B. Arthur, *Amer. Econ. Review*, 84, 406, (1994).

11. J. Casti, *Complexity*, 1, no. 5, 7, (1995/96); M. A. R. de Cara, O. Pla, F. Guinea, to appear, *EPJ*, (1999); D. Challet and Y.-C. Zhang, *Physica A*, 246, 407, (1997) and 256, 514, (1998).

12. R. E. Lucas, *Econometrica* 46, 1429 (1978).

13. W. B. Arthur, J. H. Holland, B. LeBaron, R. Palmer, and Paul Tayler, in *The Economy as an Evolving Complex System II*, 15, (Op. Cit.); W. B. Arthur, *Complexity*, 1 (1), 20, (1995).

14. See K. Lindgren's classic paper in *Artificial Life II*, C. G. Langton, C. Taylor, J. D. Farmer, S. Rasmussen, Eds. (Addison-Wesley, Reading, Mass., 1991). H. P. Young in *Econometrica* 61, 57, 1993; L. E. Blume, in *The Economy as an Evolving Complex System II*, 425, (Op. Cit.); B. A. Huberman, N. S. Glance, *Proc. Nat. Acad. Sci.* 90, 7716, (1993).

15. R. Marimon, E. McGrattan, T. J. Sargent, *J. Econ. Dynamics and Control*, 14, 329, (1990); M. Shubik, in *The Economy as an Evolving Complex System II*, 263, (Op. Cit.); W. A. Brock, P. de Lima, in *Handbook of Statistics 12: Finance*, G. S. Maddala, H. Rao, H. Vinod, Eds. (North Holland, Amsterdam, 1995)

16. T. J. Sargent, *Bounded Rationality in Macroeconomics*, (Clarendon Press, Oxford, 1993); D. A. Lane and R. Maxfield, in *The Economy as an Evolving Complex System II*, 169, (Op. Cit.); V. M. Darley and S. A. Kauffman, in *The Economy as an Evolving Complex System II*, 45, (Op. Cit.).

17. D. C. North, in The Economy as an Evolving Complex System II, 223, (Op. Cit.).

18. Y. M. Ioannides in The Economy as an Evolving Complex System II, 129, (Op. Cit.); A. P. Kirman in The Economy as an Evolving Complex System II, 491, (Op. Cit.); L. Tesfatsion in The Economy as an Evolving Complex System II, 533, (Op. Cit.).

19. P. Bak, K. Chen, J. Scheinkman, M. Woodford in *Ricerche Economiche* 47, 3, (1993); A. Leijonhufvud, in *The Economy as an Evolving Complex System II*, 321, (Op. Cit.).

20. R. Axelrod, *Am. Pol. Sci. Rev.* 80, 1095, (1986); K. Kollman, J. H. Miller, S. E. Page, in *The Economy as an Evolving Complex System II*, 461.