Artificial Adaptive Agents in Economic Theory

By JOHN H. HOLLAND AND JOHN H. MILLER*

Economic analysis has largely avoided questions about the way in which economic agents make choices when confronted by a perpetually novel and evolving world. As a result, there are outstanding questions of great interest to economics in areas ranging from technological innovation to strategic learning in games. This is so, despite the importance of the questions, because standard tools and formal models are ill-tuned for answering such questions. However, recent advances in computer-based modeling techniques, and in the subdiscipline of artificial intelligence called machine learning, offer new possibilities. Artificial adaptive agents (AAA) can be defined and can be tested in a wide variety of artificial worlds that evolve over extended periods of time. The resulting *complex adaptive systems* can be examined both computationally and analytically, offering new ways of experimenting with and theorizing about adaptive economic agents.

Many economic systems can be classified as complex adaptive systems. Such a system is *complex* in a special sense: (*i*) It consists of a network of interacting agents (processes, elements); (*ii*) it exhibits a dynamic, aggregate behavior that emerges from the individual activities of the agents; and (*iii*) its aggregate behavior can be described without a detailed knowledge of the behavior of the individual agents. An agent in such a system is *adaptive* if it satisfies an additional pair of criteria: the actions of the agent in its environment can be assigned a value (performance, utility, payoff, fitness, or the like); and the agent behaves so as to increase this value over time. A complex adaptive system, then, is a complex system containing adaptive agents, networked so that the environment of each adaptive agent includes other agents in the system.

Complex adaptive systems usually operate far from a global optimum or attractor. Such systems exhibit many levels of aggregation, organization, and interaction, each level having its own time scale and characteristic behavior. Any given level can usually be described in terms of local niches that can be exploited by particular adaptations. The niches are various, so it is rare that any given agent can exploit all of them, as rare as finding a universal competitor in a tropical forest. Moreover, niches are continually created by new adaptations. It is because of this ongoing evolution of the niches, and the perpetual novelty that results, that the system operates far from any global attractor. Improvements are always possible and, indeed, occur regularly. The everexpanding range of technologies and products in an economy, or the everimproving strategies in a game like chess, provide familiar examples. Adaptive systems may settle down temporarily at a local optimum, where performance is good in a comparative sense, but they are usually uninteresting if they remain at that optimum for an extended period.

A theory of complex adaptive systems based on AAA makes possible the development of well-defined, yet flexible, models that exhibit emergent behavior. Such models can capture a wide range of economic phenomena precisely, even though the development of a general mathematical theory of complex adaptive systems is still in its early stages.[1] The AAA models complement current theoretical directions; they are

*Professor of Psychology and Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109; and Assistant Professor of Economics and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, respectively; both are External Professors, Santa Fe Institute, Santa Fe, NM 87501.

[1] It is important in this research to determine just where the potential for general solutions exists. There are simple models of cellular automata, for example, wherein the solutions to particular questions are computationally irreducible—the shortest way to analyze the system is to run the complete computation.

not intended as a substitute. Many of the most interesting questions concern points of overlap between AAA models and classical theory. As a minimal requirement, wherever the new approach overlaps classical theory, it must include verified results of that theory in a way reminiscent of the way in which the formalism of general relativity includes the powerful results of classical physics.

I. Why Study Artificial Adaptive Agents?

The AAA models have several characteristics that are not available in traditional modeling techniques. Models based on pure linguistic descriptions, while infinitely flexible, often fail to be logically consistent. Mathematical models lose flexibility, but gain a consistent structure and general solution techniques. The AAA models, specified in a computer language, retain much of the flexibility of pure linguistic models, while having precision and consistency enforced by the language. The resulting models are dynamic and are "executable" in the sense that the unfolding behavior of the model can be observed step by step. This makes it possible to check the plausibility of the behavior implied by the assumptions of the model. The precision of the definitions also opens AAA models to mathematical analysis. The ability to explore a wide range of phenomena involving learning and adaptation, linked with the rigor imposed by a computer language, provides a powerful modeling technique.[2]

The AAA models offer a way of approaching one of the major questions of present theory. Current theoretical constructs, based on optimization principles, often require technically demanding derivations. It is an obvious criticism of these constructs that real agents lack the behavioral sophistication necessary to derive the proposed solutions. This dilemma is resolved if it is postulated that adaptive mechanisms, driven by market forces, lead the agents to act as if they were optimizing (see, for example, Milton Friedman, 1953). AAA explicitly model this link between adaptation and market forces, and can thus be used to analyze the conditions under which optimization behavior will (not) occur.

Insofar as human behavior is driven by adaption, an understanding of AAA may prove to be a useful benchmark for, and provide insights into, existing human experiments (see, for example, J. A. Andreoni and Miller, 1990; Brian Arthur, 1990).[3] An experiment consisting of artificial agents allows the utility, risk aversion, information, knowledge, expectations, and learning of each subject to be carefully controlled. Moreover, at any point in the experiment, the knowledge and learning of the artificial agents can be "reset" to any desired previous state, and subtle variations of the environment can be analyzed. The strategy (as well as the behavior) of the AAA can always be explicitly analyzed, something not usually possible with human subjects. Finally, the infinite patience and low motivational needs of AAA "subjects" implies that large-scale experiments can be conducted at a relatively low cost.

A major feature of AAA models is their ability to produce emergent behavior. A wide variety of behaviors can arise endogenously, even though these behaviors, as with any model, are constrained by the initial structure. The possibilities are so rich that it is often difficult to predict on a priori grounds what behaviors and structures will emerge. It thus becomes possible to explore realms that were unanticipated when the model was defined. Analysis of these emergent phenomena should offer both insights and suggestions for new theorems about the effects of adaptive agents in economic systems.

The AAA models may also prove useful in studying economic systems that have either an absence or a plethora of theoretical solutions. Many important economic prob-

---

[2] Programming even a simple market is instructive on the limitations of both the pure linguistic and mathematical approaches.

[3] Artificial agents could also be used as "subjects" in pilot studies to identify potentially interesting new human experiments.

lems, such as double-auction strategies, multisectoral general equilibrium models, and the like, have no easily derived analytic solutions. Several AAA techniques were originally designed as optimization methods for environments that are nonlinear, noisy, discontinuous, or involve enormous search spaces. As a result, they offer useful numerical techniques for such problems in economics. At the opposite extreme are systems with multiple solutions. For example, in repeated games, the Folk theorem often admits a vast number of potential solutions. In these cases, the interaction of the adaptive systems with the economic environment may narrow the set of potential solutions. Different equilibria may have different degrees of *adaptive complexity*.

Beyond complementing current theoretical and empirical work, AAA offer the potential for unique extensions of current theory. The mechanisms generating the global behavior of a complex adaptive system can be directly observed when the computer is an integral part of the theory. For such theories, the computer plays a role similar to the role the microscope plays for biology: It opens up new classes of questions and phenomena for investigation. Problems that prove difficult for traditional mathematical approaches are often easily implemented as an AAA system. In that form, they can be dissected and modified with ease, providing new opportunities for theory generation and testing. More generally, the potential for the development of a general calculus of "adaptive mechanics" exists. A calculus of these systems would combine the advantages of analytic perspicacity with computer-driven hypothesis testing.

II. Some Current Artificial Adaptive Agent Techniques

A wide range of computer-based adaptive algorithms exist for exploring AAA systems, including classifier systems, genetic algorithms, neural networks, and reinforcement learning mechanisms. The multiplicity of techniques presents a problem for analysis. How sensitive are the results to a particular incarnation of the adaptive agent? This problem, of course, confronts any attempt to lessen the rationality postulates traditionally used in economic theory. Usually, there is only one way to be fully rational, but there are many ways to be less rational. It is important in building a theory based on AAA to construct agents that exhibit robust behavior across algorithmic choices. Current economic studies of adaptive agents rely on genetic algorithms (R. M. Axelrod, 1987; Miller, 1989; Andreoni-Miller) and classifier systems (R. Marimon et al., 1990; Arthur).

Genetic algorithms (GAs) were developed by Holland (1975) as a way of studying adaptation, optimization, and learning. They are modeled on the processes of evolutionary genetics. A basic GA manipulates a set of structures, called a *population*. Structures are usually coded as strings of characters drawn from some finite alphabet (often binary). For example, in a game context, a string might be interpreted either as a simple strategy (a rule table) or as a computer program for playing the game (a finite automaton). Depending upon the model, an agent may be represented by a single string, or it may consist of a set of strings corresponding to a range of potential behaviors. For example, a string that determines an oligopolist's production decision could either represent a single firm operating in a population of other firms, or it could represent one of many possible decision rules for a given firm. Whatever the interpretation, each string is assigned a measure of performance, called its *fitness*, based on the performance of the corresponding structure in its environment. The GA manipulates this population in order to produce a new population that is better adapted to the environment.

In execution, a GA first makes copies of strings in the population in proportion to their observed performance, fitter strings being more likely to produce copies. As a result, fitter strings are more likely to contribute to the new population. After the copies are produced, they are modified by the application of genetic operators. The genetic operators provide for the introduction of new strings (structures) that still

retain some of the characteristics of the fitter strings in the parent population.

The primary genetic operator for a GA is the *crossover* operator. The crossover operator is executed in three steps: 1) a pair of strings is chosen from the set of copies; 2) the strings are placed side by side and a point is randomly chosen somewhere along the length of the strings; 3) the segments to the left of the point are exchanged between the strings. For example, crossover of 111000 and 010101 after the second position produces the offspring strings 011000 and 110101. Crossover, working with reproduction according to performance, turns out to be a powerful way of biasing the system toward certain patterns, *building blocks*, that are consistently associated with above-average performance.

It can be proved (see Holland, 1975) that GAs are a powerful technique for locating improvements in complicated high-dimensional spaces. They exploit the mutual information inherent in the population, rather than simply trying to exploit the best individual in the population. We can liken each of the potential building blocks to one arm of an n-armed bandit. Under this interpretation, each successive generation samples the building blocks in a way that closely corresponds to the optimal solution of an n-armed bandit problem. The GA learns by biasing the search toward combinations of above-average building blocks. Reproduction and crossover are very simple operations that impose low-information and processing requirements on the agents employing them.

A *classifier system* (CS) (Holland et al., 1986) is an adaptive rule-based system that models its environment by activating appropriate clusters of rules. It uses a GA to revise its rules. Each rule is in condition/action form, and many rules can be active simultaneously. The action part of a rule specifies a message that is to be posted when the rule is activated. The condition part of a rule specifies messages that must be present for it to be activated. Thus, each rule is a simple message-processing device that emits a specific message when certain other messages are present. Overt actions affecting the environment are the result of messages directed to the system's output devices (effectors), while information from the environment is received via messages generated by its input devices (detectors). The overall system is computationally complete in the sense that any program written in a programming language, such as FORTRAN, can also be implemented by a CS.

A CS-rule does not automatically post its message when its condition part is satisfied. Rather, it enters a competition with other rules having satisfied conditions. The outcome of this competition is based on a quantity, called *strength*, assigned to each rule. A rule's strength measures its past usefulness, and it is modified over time by one of the system's learning algorithms (see below). There may be more than one winner of the competition at any given time—hence a cluster of rules can react to external situations. A CS operates on large numbers of rules, with a small number of simple, domain-independent mechanisms. It provides emergent, learned capabilities for reacting to its environment.

A CS adapts or learns through the application of two well-defined machine-learning algorithms. The first algorithm, called a *bucket-brigade algorithm*, adjusts rule strengths. Each rule is treated as an intermediate producer in a complex economy, buying input messages and selling output messages. When a satisfied rule R succeeds in the competition to post its own message, it pays the rule(s) that supplied the messages satisfying its condition part. This amount is subtracted from R's strength. On the next time-step, if other rules are satisfied by R's message, and win the competition in turn, then R receives the rules' payment. R's strength is increased accordingly. The net effect of the two transactions is R's profit (loss). Some rules also act directly on the environment in a way that produces direct payoff from the environment to the system. Their strength is increased in proportion to that payoff. A rule's strength will increase over time only if it earns a profit, on average, in these transactions. Generally this happens only if the rule directly produces payoff, or else belongs to one or more causal chains leading to payoff. Under appropriate conditions, the

strengths assigned by the bucket-brigade algorithm do converge to a useful measure of the rule's contributions to system performance (Holland et al.).

In order to generate and test new approaches to the environment, the CS needs a second learning algorithm, a *rule discovery algorithm*. A GA can be used for this purpose, because the rules of a CS can be represented by strings in an appropriate alphabet, and a rule's strength amounts to a measure of its performance. The GA, by forming new rules in terms of tested, above-average building blocks, transfers experience from the past to new situations. Plausible new rules result—rules to be tested and retained or discarded on the basis of their ability to enhance the performance of the CS.

Under the combined effects of the bucket-brigade and genetic algorithms, rules become coupled in complex networks. Clusters and hierarchies of rules emerge. Over time, these substructures serve as building blocks for still more complex substructures. A CS agent can: 1) generate broad categories for describing its environment (so that experience can be brought to bear on novel situations); 2) progressively refine and elaborate the relation between categories (using experience to make distinctions and associations not previously possible); 3) use these categories to build internal models that supply the agent with expectations about the world; 4) treat all internal models as provisional (subject to confirmation or refutation as experience accumulates); and 5) generate new hypotheses that are plausible in terms of accumulated experience. Moreover, because of the bucket-brigade algorithm, these activities can proceed in an environment where payoff is intermittent or rare. Such capacities enable a CS agent that is not omniscient to act with increasing rationality.

III. Towards a Mathematics of Complex Adaptive Systems

A mathematical calculus appropriate to the study of complex adaptive systems must meet distinctive requirements. The usual mathematical tools, exploiting linearity,

fixed points, and convergence, provide only an entering wedge. In addition we need a mathematics that works in close conjunction with computer modeling techniques—one that puts more emphasis on combinatorics and algorithms. We require techniques that emphasize the emergence of structure, particularly internal models, through the generation, combination, and interaction of building blocks. The present situation seems quite similar to that of evolutionary theory prior to the development of a mathematical theory of genetic selection (R. A. Fisher, 1930).

Though there is nothing like an overall theory, there are some extant pieces of mathematics that are relevant. The schema theorems for genetic algorithms (Holland, 1975) offer some insight into processes that discover and recombine building blocks. It appears that schema theorems are special cases of a much more general formulation of the effects of recombination in evolution. This formulation should bring some of the more sophisticated tools of mathematical genetics to bear on adaptive agent models. Mathematical work aimed at understanding the evolution of CS may also be useful. The progressive development of hierarchical organization can be treated as the addition of levels to a quasi homomorphism (Holland et al.).

Perpetual novelty can be modeled by a regular Markov process in which each of the states has a recurrence time that is large with respect to any feasible observation time. Equivalence classes can be imposed and used as the states of a *derived* Markov process (Holland, 1986). Work by Miller and S. Forrest (1989), based on S. A. Kauffman's (1984) studies of random graphs, provides additional insights into the emergent structures of CSs.

IV. Conclusions

The AAA research complements ongoing theoretical and empirical work, allowing exploration and analysis of previously inaccessible phenomena. What are the future prospects for this line of inquiry? Early work with AAA in economics has shown that they can acquire sophisticated behavioral patterns. Observation of the course of learning

in these AAA has already increased our understanding of some economic issues. Even limited AAA open up new avenues for analyzing decentralized, adaptive, and emergent systems. Steady advances in computation and AAA modeling offer ever more powerful tools for programming artificial worlds. By executing these models on a computer we gain a double advantage: (*i*) An experimental format allowing free exploration of system dynamics, with complete control of all conditions; and (*ii*) an opportunity to check the various unfolding behaviors for plausibility, a kind of "reality check." Whether or not agents in such worlds behave in an optimal manner, the very act of contemplating such systems will lead to important questions and answers.

REFERENCES

Andreoni, J. A. and Miller, J. H., "Auctions with Adaptive Artificially Intelligent Agents," Santa Fe Institute Working Paper, No. 90-01-004, 1990.

Arthur, W. B., "A Learning Algorithm that Replicates Human Learning," Santa Fe Institute Working Paper, No. 90-026, 1990.

Axelrod, R. M., "The Evolution of Strategies in the Iterated Prisoner's Dilemma," in L. D. Davis, ed., *Genetic Algorithms and Simulated Annealing*, Los Altos: Morgan-Kaufmann, 1987, 32–41.

Fisher, R. A., *The Genetical Theory of Natural Selection*, Oxford: Clarendon Press, 1930.

Friedman, M., *Essays in Positive Economics*, Chicago: University of Chicago Press, 1953.

Holland, J. H., *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press, 1975.

_____, "A Mathematical Framework for Studying Learning in Classifier Systems," in J. D. Farmer et al., eds. *Evolution, Games and Learning*, Amsterdam: North-Holland, 1986, 307–17.

_____et al., *Induction: Processes of Inference, Learning, and Discovery*, Cambridge: MIT Press, 1986.

Kauffman, S. A., "Emergent Properties in Randomly Complex Automata," *Physica D*, 1984, *10*, 145–56.

Marimon, R., McGratten, E. and Sargent, T. J., "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents," *Journal of Economic Dynamics and Control*, May 1990, *14*, 329–73.

Miller, J. H., "The Coevolution of Automata in the Repeated Prisoner's Dilemma," Santa Fe Institute Working Paper, No. 89-003, 1989.

_____ **and Forrest, S,** "The Dynamical Behavior of Classifier Systems," in J. D. Schaffer, ed., *Proceedings of the Third International Conference on Genetic Algorithms*. San Mateo: Morgan-Kaufmann, 1989, 304–10.