Brief Papers

Inductive Reasoning and Bounded Rationality Reconsidered

David B. Fogel, Kumar Chellapilla, and Peter J. Angeline

Abstract— **Complex adaptive systems have historically been studied using simplifications that mandate deterministic interactions between agents or instead treat their interactions only with regard to their statistical expectation. This has led to an anticipation, even in the case of agents employing inductive reasoning in light of limited information, that such systems may have equilibria that can be predicted *a priori*. This hypothesis is tested here using a simulation of a simple market economy in which each agent's behavior is based on the result of an iterative evolutionary process of variation and selection applied to competing internal models of its environment. The results indicate no tendency for convergence to stability or a long-term equilibrium and highlight fundamental differences between deterministic and stochastic models of complex adaptive systems.**

Index Terms—**Bounded rationality, complex adaptive systems, El Farol problem, inductive reasoning.**

I. Introduction

By their very nature, complex adaptive systems are difficult to analyze and their behavior is difficult to predict. These systems, which include ecologies and economies, involve a population of purpose-driven agents, each acting to obtain required resources in an environment. The conditions these agents face vary in time both as a consequence of external disturbances (e.g., weather) and internal cooperative and competitive dynamics. Moreover, such systems are often extinctive, where those agents that consistently fail to acquire necessary goods (e.g., food, shelter, monetary capital) are eliminated from the population. The essential mechanisms that govern the dynamics of complex adaptive systems are evolutionary: random variation of agents' behavior coupled with selection in light of a nonlinear, possibly chaotic, environment. By consequence, reductionist, linear piecemeal dissection of complex adaptive systems rarely provides significant insight. The behavior of each agent is almost always more than can be assembled from the "sum of its parts" and interactions with its predators and prey, its enemies and allies. The fabric of these complex systems is tightly woven, and no examination of single threads of the fabric in isolation, no matter how exacting, can provide a sufficient understanding of the integrated tapestry.

Manuscript received June 30, 1998; revised November 19, 1998 and March 4, 1999.

D. B. Fogel is with Natural Selection, Inc., La Jolla, CA 92037 USA (dfogel@natural-selection.com).

K. Chellapilla is with the Department of Electrical and Computer Engineering, University of California at San Diego, La Jolla, CA 92093 USA (kchellap@ece.ucsd.edu).

P. J. Angeline is with Natural Selection, Inc., Vestal, NY 13850 USA (angeline@natural-selection.com).

Publisher Item Identifier S 1089-778X(99)04553-1.

Whereas the reductionism of traditional analysis fails to treat the holistic qualities of complex adaptive systems, computer simulations have been used to model low-level interactions between agents explicitly [1], [2]. Given such a model, attention is then focused on its emergent properties, patterns of observed behavior that could not be predicted easily from a linear analysis of agents' interactions. It is hoped that intricate computer simulations will provide useful tools for accurately forecasting the behavior of systems governed by the interactions of hundreds, or possibly thousands, of purposive agents acting to achieve goals in chaotic, dynamic environments [3].

One such environment that has received considerable recent attention is the market economy [4], [5]. The traditional view of human behavior as completely rational has given way to an alternative perspective of bounded rationality [6]. It is recognized that economic decision making, like most human judgment, is made in the face of incomplete knowledge both of the extrinsic market conditions and the expected actions of other entities. The cascade of suppositions about how other actors in the environment will react to current and projected circumstances might be best described as an "arms race of uncertainty." Rather than expect a stable outcome in the face of perfect information that is globally available to all agents, where each reasons correctly that there is only one best allocation of resources and each allocation is obvious to all involved, the more likely outcome for such market dynamics would seem to be characterized by chaotic transgressions and instability.

Surprisingly, one simulation of inductive reasoning and bounded rationality that has gained recent attention evidenced no such chaotic behavior [7]. Instead, the "economy" varied consistently around a stable point, and it was conjectured that aggregate behaviors in complex adaptive systems serve to bring about "natural attractors" where these systems will tend to return to a stable point when disturbed from equilibrium. If true, this would be a remarkable insight because it would imply that it may be possible to leverage traditional analytic tools of evolutionary stable systems [8] and game theory [9] to determine the equilibrium conditions of complex adaptive systems. Unfortunately, as the current experiments indicate, this is unlikely to be the case.

II. Background

The economic system under investigation is an idealized model of agents who must predict whether or not to commit a resource in light of the likely commitment of other agents in the environment. Commitment is time dependent, iterated over a series of interactions between the agents where previous behavior can affect future decisions. To distill the essential aspects of the model, Arthur [7] suggested the following setting based on a bar, the El Farol in Santa Fe, NM, which offers Irish music on Thursday nights.

Let each of N Irish music aficionados choose independently whether or not they will go to the El Farol on a certain Thursday night. Further, suppose that each attendee will enjoy the evening if no more than a certain percentage of the population N are present, otherwise the bar is overcrowded. To make considerations specific, let $N = 100$ and let the maximum number of people in the bar before becoming overcrowded be 60. Each agent interested in attending cannot collude with others to determine or estimate the density in the bar *a priori*; instead, they must predict how busy the bar will be based on previous attendance. Presume that data on prior weeks' attendance are available to all N individuals. Based on these data, each person makes a prediction about the likely attendance at the bar on the coming Thursday night. If their prediction indicates fewer than 60 bargoers then they will choose to attend; otherwise they will stay home. The potential for paradoxical outcomes is clear: If everyone believes that the bar will be relatively vacant then they will attend, and instead it will be crowded; conversely, if everyone believes the bar will be crowded, it will be empty. Of interest are the dynamics of attendance over successive weeks.

Arthur [7] offered the following procedure for determining this attendance. Each individual has k predictive models and chooses whether or not to attend the bar based on the prediction offered by their current best (or active) model measured in terms of how well it fit the available weekly attendance. The active model is dependent on the historical attendance, and in turn the attendance is dependent on each individual's active model. It is evident that the class of models used for predicting the likely attendance can have an important effect on the resulting dynamics. The specifics in [7] are not clear on which models were used, but some were suggested, including 1) use the last week's attendance, 2) use an average of the last four weeks, 3) use the value from two weeks ago (a period two cycle detector), and so forth. Starting from a specified set of models assigned to each of the N individuals, the dynamics were completely deterministic. The results indicated a consistent tendency for the mean attendance over time to converge to 60. Curiously, a mixed strategy of forecasting above 60 with probability 0.4 and below 60 with probability 0.6, which would engender a mean attendance of 60 individuals, is a Nash equilibrium when the situation is viewed in terms of game theory [7]. This result implies that traditional game theory may be useful in explicating the expected outcomes of such complex systems.

But people do not reason with a fixed set of models, deterministically iterated over time. Indeed, inductive reasoning requires the introduction of potentially novel models that generalize over observed data; restricting attention to a fixed set of rules appears inadequate. A more appropriate model of the El Farol problem would therefore include both a stochastic element, where new models were created by randomly varying existing ones, and a selective process that served to eliminate models that were relatively ineffectual. Individuals would thereby improve their predictive models in a manner akin to the scientific method and evolution [10]. The results of this variant on the method of [7] are qualitatively different and do not reflect any tendency toward stability in the limit or on the average.

III. Methods

Following [7], N was set equal to 100 and the bar was considered overcrowded if attendance exceeded 60. Each individual was given $k = 10$ predictive models. For simplicity, these models were autoregressive (AR) with their output made unsigned and rounded. For the ith individual, their jth predictor's output was given by

$$\hat{x}_j^i(n) = \text{round}\left(\left|a_j^i(0) + \sum_{t=1}^{l_j^i} a_j^i(t)x(n-t)\right|\right)$$

where $x(n-t)$ was the attendance on week $(n-t)$, l_j^i was the number of lag terms in the jth predictor of individual i, $a_j^i(t)$ was the coefficient for the lag t steps in the past, and $a_j^i(0)$ represented a constant bias term. Taking the absolute value and rounding the model's output ensured nonnegative integer values. Any predictions greater than 100 were set equal to 100 (predictions greater than the total population size N were not allowed). For each individual, the number of lag terms for each of its ten models was chosen uniformly at random from the integers $\{1, \ldots, 10\}$. The corresponding lag terms (including the bias) were uniformly distributed over the continuous range $[-1, 1]$.

Prior to predicting the current week's attendance, each individual evolved its set of models for ten generations. This was somewhat arbitrary, but was chosen to allow a minimal number of iterations for improving the existing models. The evolution was conducted as follows.

1) For each individual i, one *offspring* was created from each of its $k = 10$ models (designated as *parents*). The number of lag terms in the offspring from parent j was selected with equal probability from $\{l_j^i - 1, l_j^i, l_j^i + 1\}$. If $l_j^i = 1$, then $l_j^i - 1$ was not allowed, and similarly if $l_j^i = 10$, then $l_j^i + 1$ was not allowed (the number of lags was constrained to be between one and ten, with this choice also being somewhat arbitrary but sufficient to allow considerable history to affect current predictions). The number of AR terms in each offspring thus differed by at most one from its parent. The AR coefficients of the offspring were created by adding a zero mean Gaussian random variable with standard deviation 0.1 [i.e., $N(0, 0.1)$] to each corresponding coefficient of its parent. Any newly generated AR coefficients (due to an increase in the number of lag terms) were chosen by sampling from a $N(0, 0.1)$. When this step was completed, each individual had ten parent and ten offspring AR models.
2) Each of the 20 models (ten parents and ten offspring) in every individual were evaluated based on the sum of

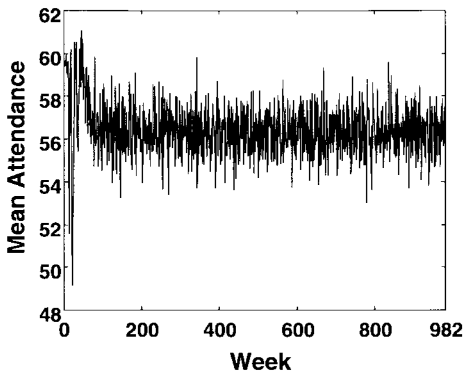

Fig. 1. The mean weekly attendance in the current simulation averaged across all 300 trials.

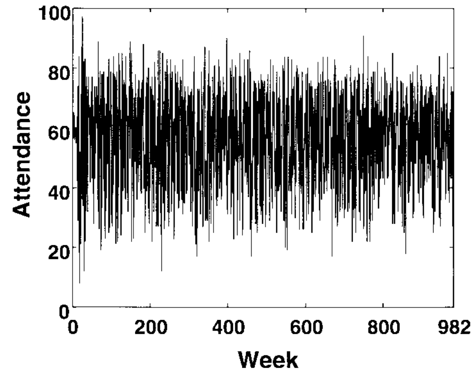

Fig. 2. The attendance observed in a typical trial.

their squared errors made in predicting the attendance of the bar during the past 12 weeks. This duration was chosen as being a sufficiently long period of time to avoid a continual transient where the population of individuals would have an insufficient sample size at each step to allow for any reasonable prediction about the current week's attendance.

3) The ten models in each individual's collection having the lowest prediction error on the past 12 weeks of data were selected to be parent models for the next generation.
4) If fewer than ten generations had been conducted, the process reverted to Step 1; otherwise each individual used their best current model (lowest error) to predict the current week's attendance. For each individual, if their prediction fell below 60 they went to the bar, otherwise they stayed home.
5) If the maximum number of weeks was exceeded, the simulation was halted, designating the completion of one trial; otherwise, the attendance for the week was recorded, the time incremented to the following week, and the process returned to Step 1.

During the first 12 weeks, attendance at the bar was initialized with truncated samples from a Gaussian random variable with mean 60 and standard deviation of five. This was meant to start the system with a sufficient sample for each individual's predictors while not biasing the mean away from the previously observed average [7] and not imposing an overwhelming variability so as to make the attendance fluctuate wildly. In all, 300 independent trials were conducted, each being executed over 982 weeks (18.83 years) so as to observe the long-term dynamics of the evolutionary system.

IV. Results

Fig. 1 shows the mean weekly attendance at the bar averaged across all 300 trials. The first 12 weeks exhibited a mean close to 59.5 resulting from the random initialization (the 0.5 decrement from 60 was an artifact caused by truncating samples to integer values). For roughly the next 50 weeks, the mean attendance exhibited large oscillations. This "transient" state transpired completely by about the one-hundredth week, with weeks 101–982 displaying more consistent statistical behavior. For notational convenience, consider this period to be described as the "steady state." The mean attendance for the steady state was 56.3155 with a standard deviation of 1.0456. This is statistically significantly different $(P < 0.01)$ from the previously observed mean attendance of 60 offered in [7]. Further, as a mean over 300 trials, the variability depicted in Fig. 1 is more than an order-of-magnitude lower than that of each single trial and the individual dynamics of each trial have been averaged out. Fig. 2 depicts the results of a typical trial having a mean steady-state attendance of 56.3931 and a standard deviation of 17.6274. None of the 300 trials showed convergence to an equilibrium behavior around the crowding limit of 60 as observed in [7], nor were any obvious cycles or trends apparent in the weekly attendance.[1] The introduction of evolutionary learning to the system of agents had a marked impact on the observed behavior: The overall result was one of chaos and large oscillations rather than stability and equilibria. Indeed, describing the dynamics of a system with behavior as shown in Fig. 2 in terms of its mean does not appear useful.

Rather than seek explanations of the stochastic system's behavior in terms of stable strategies, the essential character of the weekly attendance (i.e., the system's "state") can be captured as a simple first-order random process (higher-order effects are present because of the available time window for each agent, but as shown below, these effects are not essential to describing the behavior of the system). These stochastic models have proved useful in describing the long-term behavior of many evolutionary optimization algorithms (commonly viewed as Markov chains) [11]. Such procedures are typically designed such that only the current composition of individuals in the evolving population provides a basis for determining the next-state transition probabilities, and these probabilities are invariant for a particular population regardless of time. These characteristics would also appear to hold for the agent-based system governing attendance at the bar (with the above caveat).

Each of the first-order attendance transitions from week to week across all 300 trials were tabulated. These are shown in Fig. 3 as a normalized state transition matrix. Under the assumption that the transition probabilities are stationary and no memory of states traversed prior to the current state

[1] The attendance data from weeks 101 to 982 were examined in each trial using 1) their spectral density to test for cycles, 2) regression analysis to test for a slope significantly different from zero, and 3) a computation of their fractal dimension to test for chaotic properties. None of these tests revealed any statistically significant positive findings.

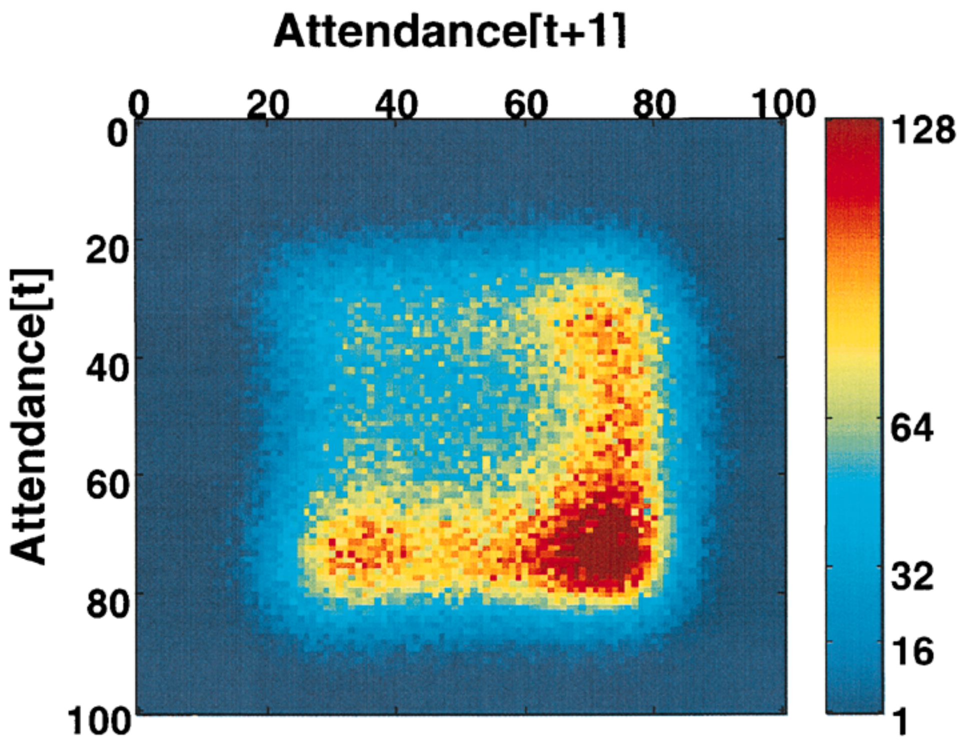

Fig. 3. The normalized one-step state transition matrix generated by tabulating all first-order transitions observed across all 300 trials from weeks 101 to 982. The color intensity reflects the normalized frequency of occurrence.

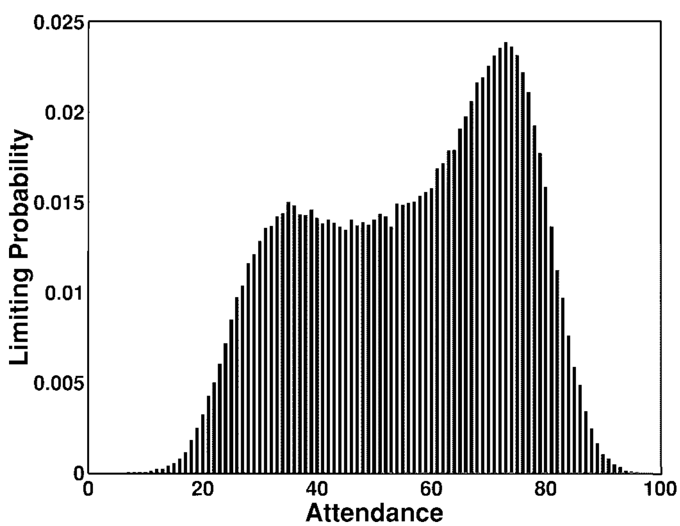

Fig. 4. The probability associated with each possible state [0–100] obtained by iterating the state transition matrix to its limiting distribution.

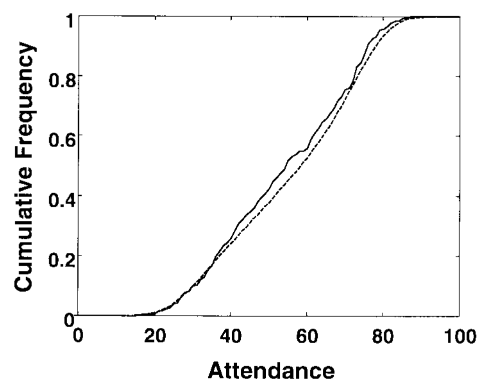

Fig. 5. The cumulative frequency of attendance in week 982 observed over 300 additional independent trials (solid) depicted against the cumulative distribution function obtained by summing the limiting probabilities of each state (dashed) as shown in Fig. 5.

is involved, the limiting probabilities of each state can be determined by raising the transition matrix to the nth power as $n \rightarrow \infty$. Beyond some value of n, the rows of the transition matrix converge to the limiting probabilities (i.e., the starting or current state is irrelevant to the long-term probabilistic behavior of the Markov chain). Fig. 4 shows the probability mass function indicating this limiting behavior, which settled to successive differences of less than 10^{-15} after 18 iterations. To provide an independent test of the hypothesis that the behavior of the stochastic agent-based system could be captured as a Markov chain, 300 additional trials were executed and each final weekly attendance at week 982 was recorded. Fig. 5 shows the cumulative frequency of these attendance figures, which appears to be in agreement with the cumulative distribution function obtained by summing the limiting probability masses for each state (see Fig. 4). A Kolmogorov–Smirnov test indicated no statistically significant difference between the proffered limiting distribution and the observed data $(P > 0.3)$; however, one assumption for the test is that the variables in question should be continuous. Thus to provide an additional examination, the observed and expected frequencies of attendance in the ranges [0–24], [25–50], [51–75], [76–100] were determined and a chi-square test again indicated no statistically significant difference between the observed and expected frequencies $(P > 0.25)$. For final corroboration, all of the weekly attendance figures for weeks 83–982 in each of the new 300 trials were tabulated. The histogram of these data appears in Fig. 6 and provides clear

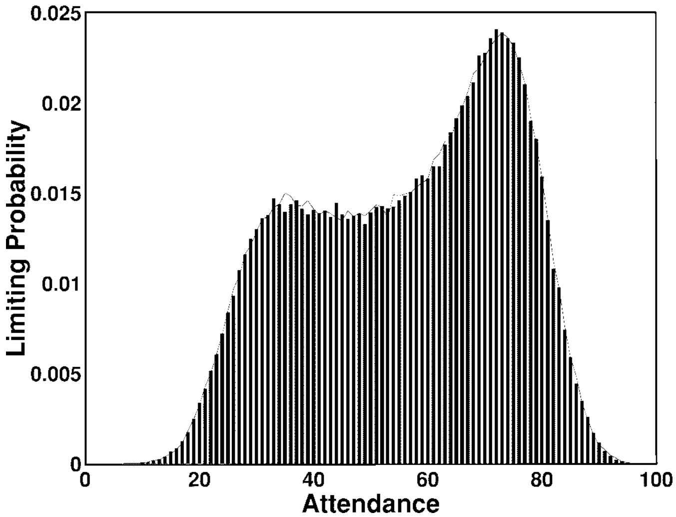

Fig. 6. The histogram of weekly attendance over weeks 83–982 observed over 300 additional independent trials. For comparison, the solid line represents the expected limiting probabilities generated by the prior 300 trials under the assumption that the system dynamics are captured as a Markov chain.

agreement with the frequencies anticipated when viewing the process as a Markov chain. Convergence to an equilibrium should not be expected from a system that is well described as a Markov chain with these limiting probabilities.

V. Discussion

The system studied here is only slightly more complex than that offered in [7]. It is certainly a highly idealized simulation of a market economy. Each agent in a constant-size population was only allowed linear predictive models with an AR form and a window into the past that was restricted to no more than three months. Moreover, the process for generating new models was a relatively simple mutation of existing coefficients and model structure. One could easily imagine variations that allowed agents to migrate to and from the city, employ generalized nonlinear predictive models, collaborate or collude with other agents, and so forth. Yet none of these more sophisticated procedures were required to generate statistically significantly different behavior from that obtained in [7].

Arthur [7] recognized the potential deficiency of mandating strictly deterministic models but suggested that any qualitative change from the previous observations would be surprising. In retrospect, perhaps there really should be no surprise. In every case of simulating complex adaptive systems, the emergent properties are strictly dependent on the "rules" preprogrammed by the investigator. Unfortunately, the results of the interactions of agents in light of even mildly complicated rules can lead to behaviors that are "surprising." This merely reflects our own ignorance, our own inability to foresee what was predestined. This inability is heightened when faced with stochastic as opposed to deterministic models. Consequently, the traditional approach in such circumstances is to either assume away the noise or average it out of consideration.

For example, in the analysis of evolutionary stable strategies (ESS's) there is often an assumption that no mutation occurs during reproduction [8], and yet this must surely be an important if not mandatory consideration in predicting the behavior of evolutionary systems. Further, in similar analyses, the outcome of encounters is often based on the mean reward to each individual. In the classic hawk-dove game, when two hawks meet, with equal probability one hawk wins and the other loses and is injured. Typical fitness values assigned are 50 and -100, respectively [13]. But the evolutionarily stable outcome, the ESS, is determined by manipulating all hawk-hawk encounters to offer a mean score of -25 to each hawk, even though neither hawk ever receives this payoff in the actual game. Models that act only on the expectation of statistical outcomes can generate altogether different behaviors than those that explicitly treat the randomness that is inherent to such circumstances [13]. When random effects are known to exist in the physical system being modeled, there must be compelling reasons for abstracting out that randomness in simulation; otherwise, the results should be viewed with caution, if not skepticism.

All models are by necessity incomplete. But there appears to be an important qualitative difference between even simple models of complex adaptive systems that rely on random variation and selection, as opposed to those that rely solely on the deterministic manipulation of fixed rules of behavior. The latter can be limited to explore only a small portion of the available space of possible strategies, while the former can be constructed such that no path is impossible [14]. The introduction of even a small degree of random variation can result in a markedly different process, stochastic in nature, with little or no qualitative agreement to the deterministic version. Models of complex adaptive systems that include agents with inductive reasoning in the face of limited information and capabilities (i.e., bounded rationality) which predict convergence to stable behavior are *ipso facto* suspect.

References

[1] C. Taylor and D. Jefferson, "Artificial life as a tool for biological inquiry," in *Artificial Life: An Overview*, C. Langton, Ed. Cambridge, MA: MIT Press, 1995, pp. 1–13.
[2] J. Carnahan, S.-G. Li, C. Costantini, Y. T. Toure, and C. E. Taylor, "Computer simulation of dispersal by *Anopehles gambiae s.l.* in West Africa," in *Artificial Life V*, C. G. Langton and K. Shimohara, Eds. Cambridge, MA: MIT Press, 1997, pp. 387–394.
[3] J. L. Casti, *Would-Be Worlds*. New York: Wiley, 1997.
[4] L. Tesfatsion, "An evolutionary trade network game with preferential partner selection," in *Evolutionary Programming V*, L. J. Fogel, P. J. Angeline, and T. Bäck, Eds. Cambridge, MA: MIT Press, 1996, pp. 45–54.
[5] N. Vriend, "Self-organization of markets: An example of a computational approach," *Comput. Econ.*, vol. 8, pp. 205–231, 1995.
[6] H. A. Simon, *Models of Bounded Rationality*. Cambridge, MA: MIT Press, 1982.
[7] W. B. Arthur, "Inductive reasoning and bounded rationality," *Amer. Econ. Assn. Papers Proc.*, vol. 84, pp. 406–411, 1994.
[8] J. Maynard Smith, *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge Univ. Press, 1982.
[9] J. Nash, "Equilibrium points in n-person games," *Proc. Nat. Acad. Sci.*, vol. 36, pp. 48–49, 1950.
[10] L. J. Fogel, A. J. Owens, and M. J. Walsh, *Artificial Intelligence Through Simulated Evolution*. New York: Wiley, 1966.
[11] G. Rudolph, "Convergence analysis of canonical genetic algorithms," *IEEE Trans. Neural Networks*, vol. 5, no. 1, pp. 96–101, 1994.
[12] R. Dawkins, *The Selfish Gene*, 2nd ed. Oxford: Oxford Univ. Press., 1989.
[13] D. B. Fogel, G. B. Fogel, and P. C. Andrews, "On the instability of evolutionary stable strategies," *BioSystems*, vol. 44, pp. 135–152, 1997.
[14] D. Hofstadter, *Fluid Concepts and Creative Analogies*. New York: Basic Books, 1995, p. 115.