programming, and bootstrap resampling methods. Forecasts are constructed by splitting up high-dimensional data sets using a tree structure to determine appropriate partitions on the variables used for forecasting. At this point the methodology is exactly that of CART. This approach differs in that a population of trees is stored, and this population is then used to draw parents for genetic recombination according to well-defined rules for crossover and mutation. A bootstrap resampling technique keeps the tree population from overfitting to specific features of a subsample of the population. This method will be compared to the more traditional tree growing and pruning techniques used in CART over some standard economic and computer-generated data sets.

IMPACT OF THE SFI APPROACH

The economics program at SFI is having an influence on the direction of economics research conducted at major universities around the world. An example is reflected in this research report by Robert Marks from the Graduate School of Business, Stanford University, and the Australian Graduate School of Management, University of New South Wales; David Midgley from the Australian Graduate School of Management, University of New South Wales; and Lee Cooper from the Anderson Graduate School of Management, University of California, Los Angeles.

Coffee Price Wars and the Genetic Algorithm: Estimating Strategies in a Real-World Oligopoly

Robert E. Marks, David Midgley, Lee Cooper

Within the study of strategic interaction, an area in economics and game theory, the recent focus has been on issues of complexity. In game theory, use has been made of machines (finite automata) to play repeated games, mainly to refine the large number of Nash equilibria possible, but also to ask questions about the complexity of response rules (or strategies) in repeated interactions. Very recently, game-theory researchers have become interested in evolutionary approaches, closer to an earlier evolutionary approach in economics, in which, roughly speaking, firms survive if they are fit (profitable) and otherwise disappear (through bankruptcy or merger or sale), without explicitly purposeful profit maximization.

Evolutionary approaches to strategic interaction are not limited to equilibrium states, and admit the possible use of simulation experiments in which game-playing machines are

not merely a theoretical abstract, but are actually used. Machine learning becomes a process in which game-playing finite automata are tested for their fitness (profitability in an oligopolistic market of a few competing sellers) as part of a process of natural selection, in which genetic algorithms are used to generate further machines for testing.

The work described here is a combination of economics, game theory, marketing, and machine learning, which uses historical data from a mature oligopoly—the retail market for ground caffeinated coffee in U.S. supermarket chains—to simulate the activities of brand managers in setting the weekly prices and other marketing instruments (coupons, advertising features, and supermarket displays) for their brands (Maxwell House, Folgers, Chock Full O Nuts, etc.).

We hope to breed artificially intelligent adaptive brand managers who may outperform their historical forebears against the other brands. In order to do this, we will use historical data to estimate a stimulus-response pattern for the other brand managers, which entails solving the problem of optimal partitioning of the space of possible histories (the particular realization of which we assume is the stimulus to which managers respond in each period).

A further issue is that of the complexity of strategies. A simple measure of the complexity of game-playing finite automata is the minimum number of states that a machine needs in order to play any strategy (where this may entail states that are never actually used in equilibrium play, but which exist to deter off-equilibrium play by opponents, so-called punishment states). But this ignores the number of possible transitions between particular states, as has been pointed out. We might expect that larger machines (which have available subtler strategies) would perform better than smaller machines in playing repeated games, despite the robustness of Tit for Tat. Future work will extend Kristian Lindgren's work on using the genetic algorithm to breed solutions in a coevolving repeated Prisoner's Dilemma with variable-length strings, which allows the number of rounds of history to become an endogenous variable in the search for fitter machines. This will allow us to explore the question of whether more complex strategies do perform better, and will also allow us to model the effects of exogenous costs of complexity (design, operation, and maintenance) on the behavior of the machines.

Emergent Structures, July, 1993

This newsletter is a publication of the Economics Research Program of the Santa Fe Institute, located at 1660 Old Pecos Trail, Suite A, Santa Fe, NM 87501.