

UCLA

UCLA Previously Published Works

Title

Breeding Competitive Strategies

Permalink

<https://escholarship.org/uc/item/4cv5c5nk>

Journal

Management Science, 43(3)

ISSN

0025-1909

Authors

Midgley, David F
Marks, Robert E
Cooper, Lee C

Publication Date

1997-03-01

DOI

10.1287/mnsc.43.3.257

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-ShareAlike License, available at <https://creativecommons.org/licenses/by-sa/4.0/>

Peer reviewed

Breeding Competitive Strategies

David F. Midgley • Robert E. Marks • Lee G. Cooper

Australian Graduate School of Management, University of New South Wales, Australia 2052

Australian Graduate School of Management, University of New South Wales, Australia 2052

Anderson Graduate School of Management, University of California at Los Angeles, Los Angeles, California 90095

We show how genetic algorithms can be used to evolve strategies in oligopolistic markets characterized by asymmetric competition. The approach is illustrated using scanner tracking data of brand actions in a real market. An asymmetric market-share model and a category-volume model are combined to represent market response to the actions of brand managers. The actions available to each artificial brand manager are constrained to four typical marketing actions of each from the historical data. Each brand's strategies evolve through simulations of repeated interactions in a virtual market, using the estimated weekly profits of each brand as measures of its fitness for the genetic algorithm. The artificial agents bred in this environment outperform the historical actions of brand managers in the real market. The implications of these findings for the study of marketing strategy are discussed.

(Competitive Strategies; Pricing; Asymmetric Market-Share Models; Evolutionary Algorithms; Repeated Games)

Introduction

We are interested in the strategic implications of asymmetric competition. Previous work (Carpenter et al. (CCHM) 1988) has estimated the Nash-equilibrium prices and advertising expenditures for asymmetric market-share models in the cases of *no competitive reaction* and *optimal competitive reaction*. There are, however, three important limitations to building marketing plans on either of these extreme scenarios.

First, such *static*, single-period strategies do not provide insight into the actions undertaken over time by manufacturers and retailers. Strategies such as ad pulsing versus continuous exposures, or every-day-low-pricing versus deep discounting are played out over time. As was called for by CCHM, it is time to investigate dynamic, multi-period strategies.

Second, major sources of asymmetries are missing from the CCHM equilibrium analysis. There are two main sources of asymmetries. Asymmetries can arise from stable, cross-competitive effects. The "price-tiers" hypothesis (Blattberg and Wisniewski 1989), for example, indicates that when national brands go on sale, they exert competitive pressure on regional brands that the

regional brands cannot counter with their own price reductions. When regional brands go on sale, they exert pressure on the economy and private-label brands that these brands cannot return. Asymmetries can also arise from temporary differences in marketing offerings. We expect that one brand on sale by itself might gain more than if it were promoted along with four other brands in the category, but such temporal *distinctiveness* of a brand's offering produces asymmetric competition in a way that is an explicit violation of the Market-Share Theory (Bell et al. 1975) and Luce's (1959) Individual Choice Theorem. Neither of these theorems allows the choice context to have any influence on choice probability, which, of course, leads to the classic counter examples by Debreu (1960). While the CCHM study incorporated measures of *distinctiveness* into their methods for reflecting asymmetric competition, their equilibrium analysis used a simpler model that did not account for this source of asymmetries.

Third, the CCHM effort studied market share, while the great swings in sales levels we observe in retail scanner data encourage us to study the strategic implications of asymmetric *sales* response. We want to investigate

multi-period strategies, when the market response is fundamentally asymmetric in both sales volume and market share.

There are major barriers to traditional avenues of investigation. Mathematical exploration is hampered because sources of asymmetry explicitly violate the global-convexity of profit functions required by normative economic models. Numerical approximations, such as traditional hill-climbing algorithms, may be computed but the complexity of the response surface makes them expensive and difficult to use. It is also difficult to incorporate more realistic competitive strategies such as "tit for tat" into traditional models—either by mathematical formulation or numerical approximation. One major alternative to mathematical or numerical exploration is multi-period simulations, such as the Axelrod (1984) or Fader and Hauser tournaments (1988). While these have the advantage of allowing strategies to be played out over time, so far they have been undertaken only with symmetric and hypothetical market-response functions. We want to use asymmetric market-response functions that characterize brand behavior in *real* markets to study the evolution of robust strategies.

Genetic algorithms (Holland 1975, Michalewicz 1994) provide one mechanism by which we can study the evolution of strategies. The next section describes genetic algorithms and how adaptable they are to the study of marketing strategy. Our major illustration is based on competition between brands in a regional market for coffee in the United States. The asymmetric market-share competition was previously modeled and mapped in this journal (Cooper 1988) and the combination of market-share dynamics and category volume effects are modeled in Cooper and Nakanishi (1988). Profits based on these published models are the central component of the fitness function that drives genetic algorithms. We use a genetic algorithm to breed artificial agents that represent the actions of brand managers. *In the tests we conduct, these agents outperform the historical actions of brand managers in this market.* Finally, we will discuss the reasons why this might be so and what can be done to extend our approach.

While we will focus on one set of modeling techniques, and one particular market, it is important to stress that the methods we propose have greater applicability. Indeed they can be used in any marketing sit-

uation where there is a good representation of the profit consequences of competitive marketing actions. This representation might be in the form of an explicit model or it might be more of a "black-box" representation (e.g., neural net). Given such a profit function, artificial agents can be formulated and genetically optimized to play multi-period dynamic games in a robust and profitable manner. Our emphasis on asymmetric market modeling in the CCHM tradition, and on a regional coffee market, simply provides one case illustration of the overall approach.

Genetic Algorithms

We model brand managers as responding to previous states of the market when formulating their marketing actions for the next period. By "state of the market" we mean the pricing and promotional actions of each competitor and their resulting profits. Each manager's competitive strategy can be represented as a set of rules, each of which maps a previous state of the market to a set of future actions (e.g., if a competitor ran a price promotion last period, do the same this period). In genetic algorithms, these sets of rules are represented as chromosome-like *binary strings*. The performance (*fitness*) of each string is judged by the profits it produces over many competitive interactions (that is, over a multi-period game). By mating the best strings, and by random mutations, new strings are created, and hence new strategies tested in parallel, as we describe below. We use machine learning to search the vast number of possible strings for robust and optimal strategies.

In essence these binary strings are *artificial agents* that represent the brand manager in repeated games (cf. Marks 1992a). Each agent responds with a unique action to each possible state of the market¹ and is therefore basing its next move on historical data. This is realistic in that human managers do not have perfect foresight of their competitors' next moves, and forecast these from the observed and remembered actions of all sell-

¹ This is a pure-strategy machine (i.e., a strategy chosen with probability 1.0). While it is feasible to employ mixed or probabilistic strategies within this approach, we chose not to for reasons of parsimony. We also do not believe that brand managers toss coins to determine their actions.

ers, including their own. If we view human managers as *boundedly rational* (Simon 1972), then we accept limitations on their memory, computing ability, or competence at pattern recognition. In our application the principal limitation we place on our artificial agents is the number of previous states of the market that can be held in memory.² If there are p players, a possible actions per round of the game, and m rounds of memory, then the number of states is $a^{m \times p}$. This number increases very rapidly. With three players, four actions, and one round of memory there are 64 states. Increasing this to two rounds of memory creates 4,096 states.

Consider the simplest game—a two-player *Prisoner's Dilemma* with one round of memory. Players have two actions, to *cooperate* (C) or *defect* (D). There are four possible states in this example (CC, DC, CD, and DD), and 16 possible mappings from past to future actions (only a few of which have obvious names). Table 1 lists the binary representation of these strategies where "0" indicates *cooperation* and "1" indicates *defection*. For example, the strategy "tit for tat" is represented by the agent (0011). Hence, we can represent our agents as strings of ones and zeroes, where there is a one-to-one mapping from states to positions on the *string*. With only two possible actions each position on the string corresponds to a single *bit*.

With more than two possible actions, we can code the actions in binary, and use as many bits per position as needed ($\log_2 a$): two bits for four actions, three bits for eight actions, etc.

In our two-player example, given that we know the profit function for the game, we play each agent against another in all possible pairings, and for a multiperiod game. Simple tabulation of the results then indicates which of the 16 strategies yields the best profit. However, as we move to more complex games, the number of possible strategies explodes, so that enumeration becomes impractical and we use a machine-learning algorithm to find the best one.

To search for better strategies, we consider the bit strings as *chromosomes* and use a simulation of *natural*

² For convenience, we have also modeled the number of actions as equal to the number of partitions of the action space perceived by the players.

Table 1 Binary Agents for a Prisoner's Dilemma—One-Round Memory

Strategy	Possible Previous States				Phantom Memory ³
	CC	DC	CD	DD	
1. Always cooperate	0	0	0	0	00
2.	0	0	0	1	
3.	0	0	1	0	
4. Nasty trigger	0	1	0	0	10
5.	1	0	0	0	
6. Tit for tat	0	0	1	1	00
7.	0	1	0	1	
8.	1	0	0	1	
9.	0	1	1	0	
10.	1	0	1	0	
11.	1	1	0	0	
12.	0	1	1	1	
13.	1	0	1	1	
14.	1	1	0	1	
15.	1	1	1	0	
16. Always defect	1	1	1	1	11

³ Phantom memory contains the strategic actions of the players the period before the game commences. That is, the initial conditions for the competitive game. Phantom memory allows us to endogenize these initial conditions. The concept was introduced by Axelrod (1987) and named by Marks (1992). For the example in the table we would need a six-bit string—where the first four bits map from states to actions and the last two are the initial conditions.

selection and reproduction with variation. The *genetic algorithm* (GA) pits the strategies against each other and *selects* the highest-scoring strategies to procreate a new generation of strategies (Holland 1975). Over repeated iterations (generations) a population of strings is evolved which exhibits even better performance.

The GA can be thought of as an optimization method that generates and evaluates a number of alternatives per iteration. The GA selects the best of these alternatives to pass on their characteristics to the next iteration. In doing so the GA evolves a population of alternative solutions that are all improving—rather than the single solution of many hill-climbing algorithms. The population of alternative solutions gives the GA its property of *explicit parallelism*. Explicit parallelism helps overcome the two problems of hill-climbing algorithms: convergence on local optima, and creeping along flat regions in objective-function space (Goldberg 1989).

The GA uses the *fitness* (brand profits in our case) of each string in the population at each generation to breed and test a new generation of strings, that may include the best individuals from the previous generation. The new generation of strings is obtained from old strings using *evolutionary operators* such as (i) reproduction of an individual according to its performance, (ii) the crossing over of genetic material of two parents, and (iii) random mutation. This process progressively biases the genetic sampling procedure toward the use of combinations of substrings associated with above-average fitness in earlier generations (i.e., strategies characterized by higher profits). GAs gain their power by searching the set of all sub-strings and identifying and exploiting the combinations that are associated with high performance.

Axelrod and Forrest have used Holland's GA to breed strategies in a two-person, repeated *Prisoner's Dilemma* game (Axelrod 1987). This study demonstrated that GAs could replace the human programmers used in the original Axelrod tournament (Axelrod 1984). Axelrod reports that the GA evolved populations whose median member resembled *Tit for Tat* and was just as successful. In some cases the GA was able to generate highly specialized adaptations to a specific population of strategies for particular situations that performed substantially better than *Tit for Tat*. There is, of course, no best winning strategy in the repeated *Prisoner's Dilemma* game: we can only speak of a best strategy given the set of strategies it is pitted against. Hence, the set of competing strategies also defines an *environmental niche*.

Marks (1989) used an environment of five of the strategies from Axelrod's second tournament. Against this environment *Always Defect* scored 343,⁴ *Always Cooperate* 406, and *Tit for Tat* 422. But after using the GA to breed a one-round memory strategy of six bits, Marks found the following strategy: 010010 (strategy 4 in Table 1). This is not *Tit for Tat* (001100), but something nastier, a trigger strategy which Defects first, and continues to do so long as the other player lets it, but once its opponent Defects, it never Defects itself again. Its score against the environment was 447, significantly better

than *Tit for Tat's*. Marks (1992b) has shown that GAs can generate new strategies in the computer without an open tournament (including for more generalized versions of the *Prisoner's Dilemma* such as that used in the Fader and Hauser (1988) tournament).

After Axelrod's pioneering study, other applications of genetic algorithms to economics have appeared (Marks 1992a, Marimon et al. 1990, Arthur 1990, and Arifovic 1994). Hurley et al. (1994) review the application of evolutionary algorithms in management science (including GAs). They find a number of applications to job shop scheduling, to financial risk, trading and portfolio management, and to some organizational problems. However, besides our early work (Marks et al. 1993) they report only one other investigation of the application of GAs to marketing problems, namely the work of Balakrishnan and Jacob (1992) on optimal product design. Hurley et al. consider that GAs have wide applicability to marketing management. They go on to provide simple illustrations of how GAs might be applied to site location and market segmentation. They note that the optimization of marketing strategies is another area of considerable potential. We believe our work is the first application in this area, and the first application in marketing which is based on empirical data.

We agree with Hurley et al. that GAs are well suited to the study of asymmetric competition. They do not require well-behaved, differentiable, globally convex objective functions. Indeed, provided we can associate a fitness with each market outcome, GAs do not require an explicit objective function at all—which provides the opportunity for an exhaustive study of patterns of strategic behavior. Our challenges are (i) to develop artificial agents that produce realistic strategies for asymmetric markets, and (ii) to combine these agents with market-response functions that translate their strategies into fitness measures. With these agents and fitness measures we can simulate market behavior for a large number of rounds—allowing us to assess the effectiveness and robustness of a strategy over a number of competitive interactions. The market simulations would occur one at a time with each artificially intelligent agent being bred against a set of rivals. This set of rivals is the niche in which competition occurs. For n brands and m alternative agents we can then use m^n niches each to

⁴ These scores are generated from the payoff function defined by Axelrod.

search for the most profitable agent for the i th brand, using the GA.

A hypothetical example may help clarify this procedure. Suppose we have a two-brand market and have chosen a population of three agents to represent possible strategies for each brand. The two brands are shown as A and B in Table 2 and the alternative agents as a_1 , a_2 , a_3 , b_1 , b_2 , and b_3 . The table shows the set of nine possible niches or competitive games that could be tested (g_1 , g_2 , etc.). Each game would be played for the desired number of rounds before computing the profit performance, and there would be three profit scores for each string. With a GA we select the two best agents (say a_1 and a_3) and "cross" their characteristics to see if we can produce a better agent in the next iteration. The nine games would then be repeated, and so on until no further improvement in profit performance is observed.

We can also take each of the m separately bred artificial agents from the final generation and separately play it against the actual history of the other $n - 1$ brands, and assess its performance against that achieved by human brand managers. That is, we can then ask if our procedure evolves a strategy for Folgers that is more profitable than Folgers brand management was. If so, then our proposed methods have value. But, before we examine a real market, we need to justify our choice of genetic algorithms and binary representations for our strategies—in preference to other evolutionary algorithms and representations.

GAs are iterative procedures that maintain a population of solutions to the focal problem, each of which is implemented as a data structure. These data structures can be binary or decimal integers, real numbers, or matrices. We use binary strings for our strategies, following the classical work of Holland (1975). Our reasons for this choice are partly convenience: we are able to use a well-accepted package—Genesis—which uses binary strings; and partly reassurance: the foundations of genetic algorithms are derived in terms of binary representation (Nix and Vose 1992, Michalewicz 1994). In common with most classical algorithms, Genesis uses fixed-length binary strings and three operators: selection, binary mutation, and crossover.

Evolutionary algorithms have also appeared with other representations, variable length strings, and modified genetic operators. Although Koza (1992) argues

Table 2 Hypothetical Niches, Two Brands and Three Agents

Brand A Agents	Brand B Agents		
	b_1	b_2	b_3
a_1	g_1	g_2	g_3
a_2	g_4	g_5	g_6
a_3	g_7	g_8	g_9

that these newer algorithms are more appropriate than string-based schemes for many problems, we believe that strings are the best choice for our mapping function. This debate centers on the degree to which one wishes to explicitly encode the rules that translate previous states to future actions. On the one hand there is the classical approach, which provides a position on the string for every possible state, and codes directly from each position to an action. On the other hand there are newer methods such as genetic programming, in which each strategy takes the form of a tree—each branch of which is labeled with an operator, and each leaf labeled with an action. If we were interested in the structure (or genotype) of the best solutions, and we had the knowledge on which to base a genetic program, then the extra effort to encode the artificial agent explicitly might be worthwhile. However, we lack detailed enough knowledge to formulate the genetic program, and we are, in any event, more interested in the emergent behavior (the phenotype) at this early stage of our research. Moreover, the theoretical basis for the newer evolutionary algorithms is only now being developed (Michalewicz 1994). It may be that these heuristic approaches will turn out to be well founded, but our preference is to rely on the better understood and more accessible classical formulation.


Asymmetric Competition in a Regional U.S. Coffee Market

Choice of Market Example

We want to work with an example of competition that exhibits four aspects of real markets.

(1) *Differential effectiveness of marketing-mix instruments across brands.* Each brand may have its own

Figure 1 Prices and Sales in Chain One



unique sensitivity to consumer response to its marketing action. The allocation of equal resources to advertising, for example, can return more to some brands than to others.

(II) *Stable cross-competitive effects.* Through positioning, historical rivalries, or other mechanisms, some subgroups of brands compete more (or less) intensely with each other than do other brands. The symmetric-market hypothesis (cf. Bell et al. 1975) dictates that gains or losses in share for one brand are distributed among others strictly in proportion to their market shares. The reality we see in markets is that some brands gain more from the losses of certain rivals than would be dictated by market share alone, while other brands are more insulated by competitive boundaries than the symmetric-market hypothesis would allow.

(III) *Asymmetries due to the temporal distinctiveness of marketing actions.* That is, we want to use methods that reflect the role of choice context on what brands are chosen. The basic thesis here is that marketing actions must be distinctive to be effective.

(IV) *The dramatic swings in volume that characterize promotion response we see in store-tracking data at a disaggregate level.* Scanner data have taught us that, when viewed at the store or chain level, market response to tactical market-mix decisions is abrupt and dramatic.

The retail coffee market modeled by Cooper and Nakanishi (1988) satisfies all four criteria. There are nine brands: Folgers, Regular Maxwell House, Maxwell House Master Blend, Hills Brothers, Chock Full O'Nuts, Chase and Sanborne, Yuban, an aggregate of premium brands called All Other Branded (AOB), and an aggregate of All Private Label (APL) brands. The data track the sales impact of price per pound (net of coupons redeemed), major newspaper ads, in-store displays, and store (but not manufacturer) coupons, for 52 weeks in three grocery chains. A considerable amount is known about modeling this market.

- The asymmetric market-share model (Cooper and Nakanishi 1988, p. 159–161) revealed differential effectiveness of price and promotion for all major brands (as well as differential base franchises depicted in the dif-

ferent brand-specific intercepts for the share model). This model is also the basis of competitive maps published in this journal (Cooper 1988).

- There were 44 significant cross-competitive effects on market share: 21 involving prices, 4 involving features, 9 involving displays, and 10 involving store coupons.

- In Cooper and Nakanishi's analysis the distinctiveness of marketing activity in each choice context was presented by an exp(z-score) model—in essence a multinomial logit model using standardized scores for the marketing instruments in each chain each week. If the action across brands on a particular marketing instrument was indistinct, all brands scored, in essence, a zero.⁵ The asymmetric market-share model has an R^2 of 0.93 and cross-validates to a hold-out sample with an R^2 of 0.85 (Cooper and Nakanishi 1988, p. 158).

- The asymmetric market-share model was combined with a category volume model (Cooper and Nakanishi 1988, p. 166). Total category volume was represented as being log-log for the price terms and log-linear for the other marketing instruments. The category volume model has an R^2 of 0.91 (Cooper and Nakanishi 1988, p. 167). Since sales are the product of market share and category volume, these combined models can reflect the dramatic swings in sales we see in scanner data.

The asymmetric market-share model and the category volume model have been combined into a single-shot market simulator called *Casper* (competitive analysis system for promotional effectiveness research, Cooper and Nakanishi 1988, p. 219–257). The Casper simulator merely takes as input prices and promotional activities of all brands, estimates market shares for all brands and sales volume for the whole category, and then calculates the profit implications for all brands. These profit estimates are thought to be roughly accurate.⁶

⁵ Even in the absence of stable cross-competitive effects, representing the role of context using exp(z-scores) guarantees that a choice model or market-response model is a non-IIA representation.

⁶ Profit margins and hence unit costs were estimated from publicly available corporate and SBU level accounting information rather than provided by the companies concerned. To the extent that these estimates are inaccurate, the validity of our results for the coffee market may be reduced.

Typical behavior of the major national coffee brands is to maintain the quality image through periods of high shelf price and no promotional activity, and then to cut their price and engage in newspaper advertising, in-store displays and coupon distribution (through both store and manufacturer coupons). The effect, not unexpectedly, is usually to increase sales and market share, and perhaps total profits in the market, depending on the costs of the promotions and the activities of other brands in the market. We have one year of weekly observations across the three retail chains operating in this two-city market, and Cooper and Nakanishi estimated market share and category volume models for the three retail chains (1988, p. 219–257), but for the sake of simplicity, we focus on the 52 weeks of data for Chain One. There is no conceptual difficulty with extending the approach to encompass the three chains, but interpreting the results would be more complex. In a first application we chose to limit complexity to make the interpretation of the results relatively straightforward. The data for Chain One exhibit all the aspects of market behavior mentioned above (differential effectiveness, stable cross-effects, asymmetries, and swings). Choosing Chain One also means we do not have to consider private labels—whose goals might be different from major brands—as these are not present. The overall patterns of prices and sales for the three major brands available in Chain One (Maxwell House Regular, Folgers, and Chock Full O'Nuts) are depicted in Figure 1. These three brands account for 77% of the market in Chain One.

Given the Chain One data, there are at least two ways we might breed artificial agents.

—*Closed-loop learning*. Breed populations of each brand against the history of the other seven over the complete 52 weeks. In this approach, each brand's agents are exposed to all of the other brands' actions and thus to a diversity of competitive strategies. But the competitors' of the focal brand do not react to the actions of its agent; they simply repeat history. There is no way around the static nature of the historical data, since it does not reveal what the contingent strategies of the competing brands might have been (given the actions of the agent).

—*Open-loop learning*. Co-evolve populations of each of the eight brands against all the other brands, using

the *Casper* model to estimate the profits generated from each 52-week game, but with all actions generated by artificial agents rather than by history. In this approach we would co-evolve eight populations each of g agents and, as discussed, this can be done by way of g^8 experiments. This is analogous to breeding the agents in a laboratory experiment rather than the field as above. The advantage of this approach is that all agents "learn" how to react to the contingent strategies of other brands—rather than the unchanging patterns of historical actions. We might then try the best artificially bred agents for each brand against the historical actions of the other seven over 52 weeks.

Two tests of the artificial agents are explicit in the second and preferred approach. One test is their profit performance against other agents in the laboratory; the other is the field test of each against the historical actions of the other brands. Neither of these tests is perfect. The laboratory test is, of course, artificial. Moreover, because of convergence of behavior, it results in a reduced set of actions being expressed, and so a smaller number of positions on each string being selected for. The field test suffers from the lack of learning noted above. But, the only better tests we can envisage are to play an artificial agent against the future actions of brand managers—either in a brand-management game or in the real market. We have not conducted such tests.

There are also problems of complexity with an eight-brand example, especially if a wide range of possible actions is allowed, and hence we have a large number of states of the game. Because of the number of states our agents would need to be more complex—so that their mappings from states to actions encode an adequate number of contingencies. What is more important, with only 52 weeks of data, we might not have an adequately rich environment in which to test a complex agent. By this we mean that some contingent strategies might not be invoked by the environment and therefore their fitness never tested. While our data has a theoretical maximum of 52 unique states, in practice it has fewer, because of weeks with close or identical actions. For these reasons we sought to simplify the problem.

Modeling the Coffee Players

We want to reduce the number of possible states for both computational and data reasons. We can do this

by reducing the number of rounds of memory, by reducing the number of actions of the players, and by reducing the number of strategic players. This implies that any economy will occur only with a cost to realism. So the question becomes, what can we do with the smallest sacrifice of realism?

We note that store tracking data typically only records store coupons and not manufacturers' coupons. This was true in our case. Store coupons are merely newspaper advertisements that focus on price discounts. Customers must clip and present the newspaper coupon to receive the discount. Often, however, stores have extra copies of the ad pages available so that all shoppers can avail themselves of these discounts. Given that this effect is partially represented by the price measure (which is net of coupons redeemed), we assume that the decision to use a store coupon promotion is simply a decision to lower price. Eliminating the actions associated with couponing is probably a minimal sacrifice.

Rather than considering price to be a continuous variable with a consequently high number of states, we consider only four price levels. The first is a high, cooperative (Pareto Optimal) price. This is like a shelf price. If all brands adopted this price (implicitly collude) in the long term, profits would be maximized. Second is the noncooperative, Nash-Cournot, price that maximizes a brand's one-shot profit regardless of the other prices. Third is the two-brand coalition price that maximizes the one-shot profits for two colluding brands when a third brand does not cooperate. And fourth is an envious price that maximizes the share of a brand's own profit in total one-shot profits. These are the four special prices that Fader and Hauser (1988) discuss for their *Generalized Prisoner's Dilemma* tournament. Here we cannot calculate such prices, since real markets do not provide a closed-form profit function for each brand. However, the coffee data do seem to support the notion of four price ranges and suggest that we can reduce the complexity of the price variable by categorization. For example, smoothed histograms of weekly prices by brand show a large peak relating to shelf price and three smaller peaks relating to favored types of promotions involving price discounts with features and display.

Given that each brand has a choice of four prices and the choices of whether to *display* or not, and to *feature* or

not, there are still 16 possible actions per week. However, in the historical data we observe that features and displays are highly correlated with low prices. Managers presumably only wish to incur the cost of these promotions when they want a significant price discount to be brought to the attention of shoppers. We therefore reduce the number of actions per brand per week to four, where each price level has an associated feature and display. Reducing the number of actions to four involves some sacrifice of realism, but possibly not as great as it might first appear. The four actions we chose for each brand are representative of the majority of actions taken by that brand in the 52 weeks. Thus we maintain the connection to the four prices in strategy tournaments while reflecting the styles of promotions seen in most store tracking data.

Four actions can be coded in two bits, considerably reducing the complexity of the problem. Indeed, four actions and five brands results in 1,024 states, and if each of these governs two bits, then individual brands can be modeled with agents of length 2,048 bits. Strings of this length are feasible from a computational point of view but there would still be questions about whether our environment was rich enough to evolve adequate agents. We therefore made a final simplification and decided to focus on the three major brands in this market (Folgers, Regular Maxwell House, and Chock Full O'Nuts). While all brands are involved in the market-share dynamics, only these brands are able to significantly expand or contract category volume with their marketing actions (Cooper and Nakanishi 1988).

We can model the market as having three major players, with the other brands as fringe players who act as nonstrategic price takers. This means there are only 64 possible states (three players, each with four possible actions) and this results in strings of 128 bits. A one-round memory game with three strategic players also requires six bits of phantom memory, resulting in 134 bit strings for strategies. Strings of 134 bits are not only easy to estimate, but the 52-week environment is adequate to evolve effective agents of this length. The brands emphasized in this simplification are by far the major players in this market. The *Casper* game that has been run with MBAs for seven years has conceptualized competition this way—employing only three brand teams. Another advantage of this simplification is that

the Fader and Hauser results are more readily compared with our simulations. However, perhaps the best defense of our simplifications is that the agents we breed can perform well in both laboratory and field.

We used a version of the machine-learning GA⁷ to simulate the actual behavior of the brands in a realistic manner. Again to reduce complexity we set up the algorithm using a single population of strings for the three brands rather than three separate populations. Our open-loop procedure did not use the historical pattern of actions, but only the payoffs (profits) as estimated by *Casper*. These were used to derive a $4 \times 4 \times 4$ payoff matrix for each of the three major brands. The four possible actions that define each face of this payoff cube were a *High* price to approximate the cooperative or collusive price, a *high* price to approximate the two-person coalition price, a *low* price to approximate the noncooperative, Nash-Cournot price, and a *Low* price to approximate the envious price. We not only had to determine four price levels but also the amounts of feature and display promotions we would associate with each level. The following two-step process was used to determine the price and promotion actions for each brand. First, we used cluster analysis techniques to identify the common patterns of price-promotion behavior for each brand. Second, we chose the four patterns that covered the spectrum of possible actions and occurred with high frequency in this market. However, because of the need to have adequate degrees of freedom for the cluster analysis and to use a richer environment of actions, we chose the four actions from the data for the three chains rather than Chain One alone. The chosen actions have great similarity to actions observed in Chain One, differing at most by one or two cents on price. See Table 3 for the marketing mix associated with each action for each strategic brand. The other five nonstrategic brands are held constant at their shelf prices, which is their most common action.

The agents for each brand participate in 50-week games, with all combinations of the agents for the other two brands. The number of weeks is fixed, which might lead to end-game strategies, but the use of one round of

⁷ We adapted GAUCSD, the U.C. San Diego version of GENESIS (Schraudolph and Grefenstette 1992).

Table 3 Possible Actions for Each Strategic Brand

Action	Price Per Pound	Feature (% Stores Featuring)	Display (% Stores Displaying)
<i>Folgers</i>			
Low (0)	\$1.87	95	69
low (1)	\$2.07	83	0
high (2)	\$2.38	0	0
High (3)	\$2.59	0	0
<i>Maxwell House</i>			
Low (0)	\$1.96	95	69
low (1)	\$2.33	83	0
high (2)	\$2.46	0	0
High (3)	\$2.53	0	0
<i>Chock Full O'Nuts</i>			
Low (0)	\$1.89	100	77
low (1)	\$2.02	100	65
high (2)	\$2.29	0	0
High (3)	\$2.45	0	0

memory eliminates these. Using the GA, we chose a population size of 25 agents (or strings). Typical population sizes in the GA literature are 25 or 50. As our goal is to evolve a small number of high-performing agents rather than a larger population with high average performance, we chose the smaller number.

Consequently, testing each generation of strings requires 8,125 50-round games (325 games per string per generation). While in principle there are 25^3 (or 15,625) niches to test, this application is symmetrical, which reduces the number of niches to $25 \times (25^2 + 25)/2$ or 25×325 or 8,125 games (lower triangular matrix including the diagonal). For every week of the game each brand has complete information on all previous actions, but not on profits. The GA was set to evolve agents for 100 generations and so the final total of 50 week games was 812,500—which is equivalent to over 40 million competitive interactions.⁸ However, because the actions of the agents are limited: we simply formed a look-up table of the profit implications of all possible states. The GA was set to software defaults of a per-bit crossover rate of 13.0 and a mutation rate of 0.0001 (a rate of 0.5 would

be a random search). Crossover serves to determine which portions of the strings are exchanged to create offspring, while mutation serves to avoid premature convergence by generating new strings. GA practitioners have justified these parameters on heuristic grounds, following extensive experimentation, as discussed by Goldberg (1989). The criteria used have been convergence speeds to known solutions, ability to identify global optima from many close local optima, and other desirable characteristics.

First Experiments—Unconstrained

The first computer experiments found convergence, with all brands using their *Low* price actions and not at a collusive high price. This finding is the result of including a model for category volume as well as market shares. If only shares were modeled, strategies would probably have converged on the collusive price. But historically most of the sales and profits in this market have occurred at *Low* prices with promotions, because of stockpiling, forward buying, and brand switching, rather than through increased consumption. Indeed, Gupta (1988), in studying consumer panel data from the same time, concluded that increased sales from coffee promotions came more from brand switching than from forward buying or stockpiling. Over the time horizon of our data we can consider coffee as a mature category with stable long-term consumption rates.

Second Experiments—Institutional Constraints

To increase realism, we added some institutional constraints. Chain One does an excellent job of maximizing long-term profits while not exhausting demand. Its policy is to promote (*Low*) only one major brand at a time for the duration of one week. We mimicked this policy by saying no agent could follow one week's *Low* with another *Low*, and only one agent per week could promote *Low*. Before we play an agent against other brands, we filter against strings that map any own-play of *Low* in the previous round to own-play of *Low* in the next round, and arbitrarily assign them poor profit figures. As a result such agents will be much less likely to pass on their characteristics to members of the next generation of agents. This is because the GA uses each string's performance to determine the likelihood of that string being a parent of members of the next generation (i.e., passing on sub-strings). Ties of two or more strings that,

⁸ On a Pyramid minicomputer 100 generations of 25 individuals has a run-time of approximately 75 minutes.

given the state of the oligopoly as a result of past actions, would simultaneously price at *Low* are broken by random choice; the loser(s) arbitrarily price at *high*. The unconstrained strings are evaluated without any adjustments to their profit performance.

These institutional constraints resulted in an interesting pattern of behavior in which brands alternated in pricing *Low*, with the other two brands pricing *low*, *high*, or *High*. However, these agents still priced *Low* and *low* too frequently—resulting in saturation of demand.

Third Experiments—Demand Saturation

To make the experiments even more realistic, we introduced time into the demand side by adding demand saturation. *Casper* is a one-shot, brand-planning simulator that does an excellent job of forecasting single-period demand. But while this market is very volatile in the short run, it is very stable in the long run. Similar constraints were added to the *Casper* game to keep MBAs from acting as if consumers would *bathe* in coffee just because the price dropped. Here the weekly total demand was pro-rated by the degree of over-saturation of the past seven weeks.⁹ To represent the observed saturation of demand, we first calculated the total sales volume per week, a function of the marketing actions of the three strategic brands and the remaining non-strategic brands. We then calculated the average total sales volume over the previous seven weeks and used this together with a figure for the historical average total sales volume to calculate the percentage degree of saturation. If this percentage was greater than 100%, the total sales volume for the latest week was reduced by the degree of saturation: in steady state this will mean total sales volume equal to the historical average. Then the profits of the brands were reduced for each of the three competing brands. This was achieved by using profits calculated from *Casper* by limiting total sales, as if the total market had shrunk, which is a consequence of demand saturation.

With institutional and demand constraints in place, two patterns of competition evolved. In some cases we got convergence to all *low* pricing. We speculate

that there are likely to be additional institutional constraints that prevent this behavior and so we do not present these results here. In other cases we got convergence to patterns of behavior generally similar to that observed historically in Chain One but with higher profits. Figure 2 shows the simulated behavior of the three strategic brands with the institutional and demand constraints.

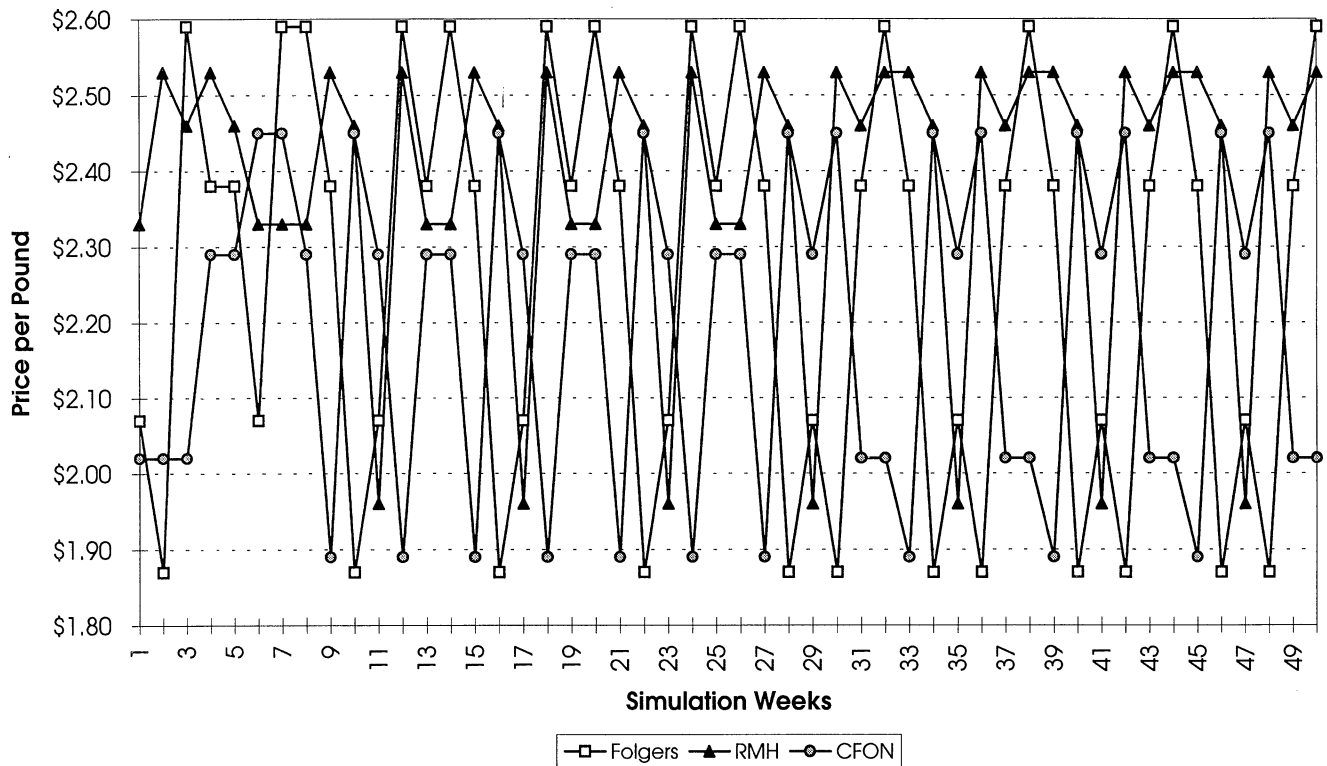
It is important to note that the results shown in Figure 2 are for three optimized agents competing against each other over 50 weeks. As such, the frequency of price competition is higher than we observe in the actual market—because the optimized agents invariably respond to the previous week's actions of their competitors. For example, the artificial agent for Folgers reduces its price 37 weeks out of 50, whereas the brand managers for Folgers only promoted 14 weeks out of 50. Similar statistics for Maxwell House are promotions on 30 weeks for the agent versus 11 weeks in the data, and for Chock Full O'Nuts 37 weeks for the agent versus 17 weeks in the data. In itself this "over-competition" is not unexpected, as our agents do not face the practical barriers met by brand managers. In our "laboratory," information on competitor actions is received instantaneously, and promotional responses can be implemented within one or two weeks (subject only to the strictures that no brand may promote *Low* on two consecutive weeks and only one brand may promote *Low* in any one week). Our artificial agents can therefore respond immediately to any competitive action, whereas human brand managers may face practical constraints. In these "laboratory" tests the agents generate profits that are very much higher than those observed in the actual market (achieving from 353% to 970% of historical averages).

Fourth Experiments—Tests Against Historical Actions

The final series of experiments is not concerned with evolving better agents; rather we took the best agents from the third series of experiments and tested each in turn against the historical actions of their seven competitors. We did this by taking an artificial agent, assigning it to one of the three strategic brands, and allowing it to choose what actions to make week-by-week over a 52-week period. However, the formulation of our agents means that they only react to the historical

⁹ The seven-week period was chosen to approximate the average interpurchase interval in this category.

Figure 2 Price Strategies for the Three Major Competitors



actions of the two other strategic competitors.¹⁰ Hence, while the historical actions of the other five nonstrategic brands are input to profit calculations, they are not “recognized” by the agent. As the GA was set to evolve a population of 25, we have 25 “best” agents, and so the test can be repeated 25 times.¹¹ Table 4 details the average profits generated by each agent over the 52 weeks—expressed as a percentage of the profits achieved by each brand’s human managers over the same period (computed by the Casper simulator from historical actions).

Table 4 shows that when the artificial agents are assigned to either the Chock Full O’Nuts or the Folgers

brand most of them do markedly better than human managers. Indeed for Chock Full O’Nuts only three agents do worse than the human managers (strings 6, 20, and 22); and for Folger’s only two agents do worse (strings 6 and 20). The best agent (string 24 in both cases) performs, respectively, 233% and 240% better than brand managers. For Maxwell House the results are not so good, with only two agents out-performing brand managers (strings 8 and 14). However, even here the best strings do produce a 20% increase in profits. The fact that the overall results are not as good for Maxwell House is amenable to two competing explanations. First, it may be the penalty for breeding one population of agents rather than separate populations for each strategic brand. If the market response to Maxwell House’s actions is sufficiently different from the other two brands’, then breeding separate populations of agents might produce better performance. Second, we note that in absolute terms Maxwell House’s profits are the highest of the three brands and may therefore be the hardest

¹⁰ In performing this test it is necessary to classify the historical actions of the other major brands into *Low*, *low*, *high*, and *High*. We did this by inspection, partitioning the price distribution into four roughly equal levels.

¹¹ The fact that we have 25 best solutions is also an example of the explicit parallelism of the GA.

Table 4 Performance of Agents Against the Actions of Brand Managers

Agent (String)	Chock Full O'Nuts	Folgers	Maxwell House
1	151%	176%	79%
2	101%	106%	72%
3	120%	135%	86%
4	114%	158%	82%
5	219%	238%	91%
6	87%	96%	76%
7	217%	174%	98%
8	200%	158%	120%
9	144%	125%	74%
10	107%	104%	77%
11	223%	226%	87%
12	187%	170%	85%
13	123%	116%	72%
14	227%	183%	116%
15	117%	112%	85%
16	111%	120%	73%
17	186%	183%	76%
18	113%	107%	76%
19	134%	108%	79%
20	89%	93%	84%
21	167%	154%	77%
22	87%	107%	75%
23	193%	148%	98%
24	233%	240%	84%
25	169%	192%	93%
Overall Mean	153%	149%	85%

The best three agents for each branch are bold.

to improve. Possibly the historical actions of Maxwell House's managers have been more effective than those of the other brands. This conclusion is supported by the fact that Maxwell House is the largest player in this market.

This test demonstrates that the "laboratory" results can be translated to the field. However, the historical test is limited because each agent is competing against "closed loop" competitors who do not learn from the actions of the artificial agent and adapt their own behavior accordingly. Therefore, we consider that what is impressive about these results is not that the agents can outperform historical actions, but that such simple agents can generate reasonable performance in this 'noisier' environment. Our agents are very simple, as they are limited to a one-round memory and only four

actions. We should ask what the patterns of behavior are that lead to such performance from these simple agents. This turns out not to be an easy question to answer because of the difficulties of presenting all the data in an understandable form. However, Figures 3 and 4 clarify the issue. Figure 3 shows the historical price actions of Folgers compared with the price actions of the "best" agent (string 24) faced with the same history of plays. Figure 4 shows the same historical actions compared with the "worst" agent (string 20).

The comparison between these two figures indicates that while the "worst" agent behaves similarly to the human manager, the "best" agent is prepared to keep the price low and promote more frequently. Although we do not present the figures here, similar conclusions can be drawn for Chock Full O'Nuts and Maxwell House.


Figure 5 illustrates the operation of the GA in this application. Three quantities are plotted in the figure: (i) the maximum profit achieved by any of the 25 agents; (ii) the average profit achieved across the 25 agents in any generation; and (iii) the percentages of agents that violate the constraints we impose.¹² These quantities are plotted against the generations of agents evolved by the GA. The brand shown is Folgers, although again the other brands have similar plots.

During the simulations of 100 generations, the best string emerges around the 45th generation, and remains unbeaten for the next 55 generations. This can be seen by the final plateau in the maximum profit series.

Figure 5 demonstrates two key properties of the GA. First, it selects those strings that can perform well in the environment they face and punishes those that cannot. Thus over the first 20 generations most of the strings that violate institutional constraints become *extinct*. The last such string becomes extinct after the 44th generation. Second, the GA progressively improves the performance of the remaining strings and reduces the variance between them. This can be seen by the fact that the plots for maximum and average performance converge

¹² These we can identify because their fitness function is set to a standard low value by our program. We cannot necessarily identify which constraint has been violated because we as yet do not fully understand the genome (bit string structures).

Figure 3 Price Strategies for Folgers—Best Agent vs. History



around the 65th generation. The result is a population of high-performing agents rather than a single best agent.


Discussion

The overall conclusion we reach is that the artificial agents price promote more frequently than do human managers. We observe the highest level of promotion when the optimized agents are competing with each other in our "laboratory." These results might be defined as the maximum competition possible in this market, with the actual market likely to show less competition. Hence, it is not surprising that placing one of these agents back into the historical market produces a lower frequency of promotion. This is the case for many of the final strings, whose behavior more resembles hu-

man managers. However, the 'best' agents still promote more frequently than do their human counterparts, and we need to speculate on the reasons for this.

One reason may be that human managers are not able to respond to competition on a week-by-week basis. More likely, they negotiate with the chains for a series of promotions to occur within a defined period (often of thirteen weeks). Major responses to competitive actions may thus occur in the next promotional period rather than by immediate adjustments to the current plan. *Institutional constraints such as promotional periods may serve to lower the frequency of promotion.* However, as yet we do not understand the impact of promotional frequency on the profits of the retailer. Our current model is of brand profit maximization, whereas in reality there is negotiated profit sharing between manufacturers and retailers. It is currently unclear to us

Figure 4 Price Strategies for Folgers—Worst Agent vs. History




whether the institutional constraints simply serve to increase chain profits and decrease brand profits, or whether they serve to jointly maximise the profits of all the parties. A future extension of our work might be just such a joint model.

Another reason human managers may promote less than our agents relates to the level of aggregation at which managers operate. Our data and models are for one chain in two cities in one state of the United States. It may be that brand and chain management structures preclude managers from managing effectively at this micro-level of detail. This is not a problem for our overall approach—which could be applied at more aggregate levels—but it does raise interesting questions as to the optimum level at which to seek to manage the promotional mix.

However, the reasons for the agents promoting more frequently than do managers may have more to do with the way we specified the agents. These reasons include the choice of one-round memory and the selection of four prices. In essence, one-round memory restricts the agent to only being directly “aware” of the most recent actions of its competitors. More rounds of memory would allow the agent to take a more balanced approach to competitive reaction, since the agent might then “assess” how aggressive a competitor’s strategy was across a greater number of instances of behavior. For example, observing that a competitor has promoted for two periods out of three implies greater aggression than if that competitor has only promoted one period out of two. In contrast, agents with one round of memory cannot distinguish these actions.

Figure 5 Evolution of String Performance for Followers



As we add rounds of memory, our artificial agents are able to make more graduated responses to competition but only at the cost of complexity—the increasing length of our strings. A three-round memory version of our agent for four actions and three players would imply 262,144 states ($4^{3 \times 3}$) and could be coded in a string of 524,288 bits plus phantom memory. Our current string has a length of 128 bits plus phantom memory, and so moving to three rounds of memory results in a string that is over 4,000 times longer. As will be discussed below, there is a less costly route to formulating better models.

The four prices we selected may have also had an impact on the behavior of the agents. While the choice of four as the number of levels has some justification in the literature (Fader and Hauser 1988), choosing the ac-

tual prices and promotional activities for each level is not straightforward. The procedure we used obviously involves assumptions and judgements, and to the extent these are incorrect so will be the behavior of our optimized agents. (It can be argued, however, that even if the actions are incorrect, the agents still perform well.) Four levels is also a simplification in that more combinations of price and promotion are observed in the market. Again, we could elaborate our models from four to eight or sixteen actions, but only at the cost of increasing the length of the string. An eight-action, three-brand agent with one round of memory implies 512 states ($8^{1 \times 3}$), and these can be coded in a string of 1,536 bits. It turns out to be less costly to increase the number of actions than to add memory. Four actions can be coded into two bits, but only three bits are needed to code eight

actions. In contrast, increasing the rounds of memory acts as a power function and therefore has a dramatic effect on the length of the strings.

Hence we have four potential explanations for the greater competitiveness of our agents as compared to brand managers, namely, the nature of institutional constraints; the effectiveness of management structures; one-round memory; and restricted price actions.

Our intuition is that the nature of institutional constraints is more likely to prove the correct explanation. However, overcoming this limitation in our agents involves a significant shift in perspective—from week-by-week actions to a pattern of actions across thirteen weeks. In essence, an action for a more sophisticated agent would be a specific promotional plan chosen from a set of possible plans. Moreover, as we shift perspective to the longer time frame, it seems likely that we will need more options for the agent to select from. This is because providing an adequate representation of a thirteen-week reality will require more alternatives than a one-week perspective.

We have recently conceived of an alternative model formulation that will potentially allow for more complex behavior to be generated from strings that are not much longer than those used here (Midgley et al. 1995). Instead of employing external cluster analysis to partition the data into discrete "actions," we might endogenize this partitioning of actions within the agent's bit string. In essence the agents would "learn" an optimal partitioning of promotional actions from previous states of the market. Since these partitions could be constructed over several periods, this alternative approach may also not require encoding several rounds of memory into the bit string. Instead the parameters of the partitioning function would be held in the bit string, and quite complex functions may only require a small number of additional bits. There is a human analogy here, since human genetic information may well set the parameters for our memory and pattern recognition algorithms. However, as yet our alternative approach is speculative. To confirm its value requires considerable further work, and a more extensive market environment than was available for the current study.

We believe we have demonstrated the potential value of our overall approach. Artificial agents can be bred which mimic the behavior of human managers and

which play dynamic, multi-period strategies. This in itself is an advance in the literature and one that has been called for (CCHM). Previous models have only provided solutions for the extreme cases of no competitive reaction or optimal (Nash-Cournot) competitive reaction. Neither of these cases is likely to be observed in actual markets, where profits are sought over a number of periods of a multi-period game. Our artificial agents seek to maximize profits over a number of periods and do so, taking account of institutional constraints, consumer response to promotions and competitive reactions. At a string length of 134 bits, these agents are actually *very* simple and are therefore capable of considerable improvement. Just as a geneticist breeds better strains of bacteria in the laboratory, we should be able to breed better agents in our simulations.

The evolutionary model is a powerful method for optimizing artificial agents to a particular competitive environment. It allows us to specify realistic agents and then to evolve those which maximise profits in this environment. However, while our overall approach is generalizable, here we have developed agents that are specific to the chosen environment. That is, the coffee agents we bred for this paper are fit to survive and prosper in one regional coffee market, but might not do well in another coffee market, and would be unlikely to survive in an entirely different product class. Equally, if the fundamental economics of coffee were to change or another major player to enter the market, our agents would not perform effectively. The evolutionary model selects the agents to match the environmental niche—which is both its strength and weakness.¹³

Fortunately, where we have an appropriate model of sales response to mix variables, and where we know the underlying cost structures of the manufacturers, we can specify an artificial agent. Applying a genetic algorithm then allows us to select agents that maximize profits over the desired time horizon. Improvements in market modeling and the wider availability of genetic algo-

¹³ Indeed the effectiveness of the agents we breed is entirely dependent on the richness of the environment they are bred in. This raises important issues of how we may best construct competitive environments in the "laboratory." In the example we have discussed, we are somewhat limited by having only one year's data for one region.

rithms will make this a practical procedure for many market situations.

What then are the managerial applications of this approach? We believe there are three. First, agents allow managers to check promotional plans against the likely response of their competitors. Promotional plans can be input for their own brand, and the competitive responses to these plans generated from the agents of the other brands. Second, agents enable managers to test "what-if" scenarios for their own brands and for those of their competitors. Both these applications may help alleviate the resistance to market modeling that is observed in some companies. Part of this resistance may stem from the static and competitively myopic nature of other current approaches. Managers expect models to simulate the consequences of actions planned over a period, and for these simulations to factor in likely competitive response. Third, agents may be useful in training junior brand managers. Agents might form the basis of games where junior managers make the decisions for one brand, and agents the decisions for competitive brands. With appropriate agents this would inject an element of realism into training by simulation games.

It is also possible that this approach has implications for regulatory agencies. We have suggested that the intensity of competition observed in the "laboratory" is a theoretical maximum. This maximum might provide a benchmark by which regulatory agencies can assess the actual intensity of competition observed in any market.

In addition to the extensions of our work discussed in this paper, there are many potential applications of GAs to marketing. GAs provide efficient and robust optimization methods for situations where the objective function is complex, nonlinear, and analytically intractable. Their explicit parallelism makes them less susceptible to convergence on local optima, and relatively effective at identifying complex patterns. As a consequence, they allow us to contemplate more complex objectives than we might otherwise, including hybrid functions, incorporating algebraic formulae and decision rules. Since the optimization of complex functions lies at the heart of many marketing applications, we can envisage applying GAs to media and sales force allocation, to modeling market and competitive response, and to optimizing profitability under various strategic

scenarios. Less obvious applications include our (unpublished) work on the use of a GA to identify consumer choice models from scanner data, and the work of Terano et al. (1995) on using a GA both to identify consumer decision rules from questionnaire data and to generate potential new products from these rules.

Of course, all of this rests on breeding more realistic agents. Our work encourages us to think that this can be done and that second-generation agents will produce even more realistic behavior. This will be achieved by conceptualizing actions over multi-period planning horizons and by endogenizing partitioning. These second-generation agents will have longer chromosomes—but to some degree this only costs us computing cycles. The more challenging problem is to develop environments that are sufficiently diverse that more complex agents can be bred to be effective and robust competitors.¹⁴

¹⁴ The authors wish to thank David Fogel for his advice, Information Resources, Inc. for providing the data used in this study and Joel Steckel for his assistance in estimating the profit margins for the brands. We also thank the editors and referees for their helpful suggestions throughout the review process. This research is supported, in part, by the Marketing Studies Center of the Anderson School at UCLA, the Australian Graduate School of Management, UNSW, the Australian Research Council, the Graduate School of Business at Stanford University, and the Santa Fe Institute. Earlier versions of this paper were presented at the Santa Fe Institute, the 1993 TIMS Marketing Science Conference in St. Louis, Missouri, the 1993 Society of Economic Dynamics and Control Conference in Nafplion, Greece, and the 1994 Australasian meeting of the Econometrics Society in Armidale, N.S.W.

References

- Arifovic, J., "Genetic Algorithm Learning and the Cobweb Model," *J. Economic Dynamics and Control*, 18 (1994), 3–28.
- Arthur, W. B., "A Learning Algorithm That Mimics Human Learning," Working Paper 90-026, SFI Economics Research Program, Santa Fe Institute, Santa Fe, NM, 1990.
- Axelrod, R., *The Evolution of Cooperation*, Basic Books, NY, 1984.
- , "The Evolution of Strategies in the Iterated Prisoner's Dilemma," *Genetic Algorithms & Simulated Annealing*, L. Davis (Ed.), Pittman, London, 1987, p. 32–41.
- Balakrishnan, P. V. S. and V. S. Jacob, "A Genetic Algorithm Based Decision Support System for Optimal Product Design," *ORSA/TIMS Marketing Science Conference*, London, United Kingdom, July 1992.
- Blattberg, R. C. and K. J. Wisniewski, "Price-induced Patterns of Competition," *Marketing Sci.*, 8, 4 (1989), 291–309.
- Bell, D. E., R. L. Keeney, and J. D. C. Little, "A Market Share Theorem," *J. Marketing Res.*, 12 (1975), 136–141.

- Carpenter, G. S., L. G. Cooper, D. M. Hanssens, and D. F. Midgley, "Modeling Asymmetric Competition," *Marketing Sci.*, 7, 4 (1988), 393-412.
- Cooper, L. G., "Competitive Maps: The Structure Underlying Asymmetric Cross Elasticities," *Management Sci.*, 34, 6 (1988), 707-723.
- and M. Nakanishi, "Standardizing Variables in Multiplicative Choice Models," *J. Consumer Res.*, 10 (1983), 96-108.
- and —, *Market-Share Analysis: Evaluating Competitive Marketing Effectiveness*, Kluwer, Boston, 1988.
- Debreu, G., "Review of R. D. Luce's *Individual Choice Behavior: A Theoretical Analysis*," *American Economic Review*, 50, 1 (1960), 186-188.
- Fader, P. S. and J. R. Hauser, "Implicit Coalitions in a Generalised Prisoner's Dilemma," *J. Conflict Resolution*, 32 (1988), 553-582.
- Goldberg, D. E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA, 1989.
- Gupta, S., "Impact of Sales Promotions on When, What, and How Much To Buy," *J. Marketing Res.*, 25 (1988), 342-355.
- Holland, J. H., *Adaptation in Natural and Artificial Systems*, Second Edition, University of Michigan Press, Ann Arbor, MI, 1992.
- Hurley, S., L. Moutinho, and N. M. Stephens, "Applying Genetic Algorithms to Problems in Marketing," *Proc. 23rd EMAC Conf., Marketing: Its Dynamics and Challenges*, J. Bloemer, J. Lemmink, H. Kasper (Eds.), European Marketing Academy, Maastricht, 1994.
- Koza, J. R., *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press, Cambridge, MA, 1992.
- Luce, R. D., *Individual Choice Behavior*, Wiley, New York, 1959.
- Marimon, R., E. McGrattan, and T. J. Sargent, "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents," *J. Economic Dynamics and Control*, 14 (1990), 329-373.
- Marks, R. E., "Niche Strategies: The Prisoner's Dilemma Computer Tournaments Revisited," AGSM Working Paper 89-009, 1989.
- , "Repeated Games and Finite Automata," *Recent Advances in Game Theory*, J. Creedy, J. Eichberger, and J. Borland (Eds.), Edward Arnold, London, 1992a.
- , "Breeding Optimal Strategies: Optimal Behaviour for Oligopolists," *J. Evolutionary Economics*, 2 (1992b), 17-38.
- , L. G. Cooper, and D. F. Midgley, "Breeding Optimal Strategies for Multiple Period Competitive Games in Mature Markets," *ORSA/TIMS Marketing Science Conference*, St. Louis, MO, March 1993.
- Michalewicz, Z., *Genetic Algorithms + Data Structures = Evolutionary Programs*, 2nd edition, Springer Verlag, Berlin, 1994.
- Midgley, D. F., R. E. Marks, and L. G. Cooper, "How Boundedly Rational Are Brand Managers?," *INFORMS Marketing Science Conference*, Sydney, Australia, July 1995.
- Nakanishi, M., L. G. Cooper, and H. H. Kassarian, "Voting for a Political Candidate Under Conditions of Minimal Information," *J. Consumer Res.*, 1 (1974), 36-43.
- Nix, A. E. and M. D. Vose, "Modeling Genetic Algorithms with Markov Chains," *Ann. Math. and Artificial Intelligence*, 1992.
- Schraudolph, N. N. and J. J. Grefenstette, A User's Guide to GAUCSD 1.4, Technical Report CS92-249, UCSD CSE Department, La Jolla, CA, 1992.
- Simon, H. A., "Theories of Bounded Rationality," in C. B. McGuire and R. Radner (Eds.), *Decision and Organisation: A Volume in Honor of Jacob Marschak*, North-Holland, Amsterdam, 1972.
- Terano, T., Y. Ishino, and K. Yoshinaga, "Integrating Machine Learning and Simulated Breeding Techniques to Analyze the Characteristics of Consumer Goods," in J. Biethahn and V. Nissen (Eds.), *Evolutionary Algorithms in Management Applications*, Springer, Berlin, 1995.

Accepted by John C. Henderson; received December 15, 1994. This paper has been with the authors 6 months for 2 revisions.