Evolutionary Algorithms in Management Applications, ed. by J. Biethahn and V. Nissen, (Berlin: Springer-Verlag), 1995.

Adaptive Behaviour in an Oligopoly

Robert E. Marks[1], David F. Midgley[1], and Lee G. Cooper[2]

[1] Australian Graduate School of Management, The University of New South Wales, Sydney, NSW 2052, AUSTRALIA

[2] Anderson Graduate School of Management, University of California, Los Angeles, CA 90024-1481, USA

Abstract. Advances in game theory have provided an impetus for renewed investigation of the strategic behaviour of oligopolists as players in repeated games. Marketing databases provide a rich source of historical evidence of such behaviour. This paper uses such data to examine how players in iterated oligopolies respond to their rivals' behaviour, and uses machine learning to derive improved contingent strategies for such markets, in order to provide insights into the evolution of such markets and the patterns of behaviour observed. The paper is an application of repeated games and machine learning to adaptive behaviour over time in the retail market for ground coffee. Using empirical data on the weekly prices and promotional instruments of the four largest and several smaller coffee sellers in a regional U.S. retail market, and using a market model to predict sellers' market shares and profits in response to others' actions in any week, we examine the adaptive strategic behaviour of the three largest sellers. We model the sellers' strategic behaviours as finite automata with memory of previous weeks' actions, and use the Axelrod/Forrest representation of the action function, mapping state to action. We use a genetic algorithm (GA) to derive automata which are fit, given their environment, as described by their rivals' actions in the past and the implicit demand for coffee.

Keywords. iterated synchronised oligopoly, asymmetrical competition, pricing, marketing strategies, stimulus-response behaviour

1 The Issues

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

First, such *static*, single-period strategies do not provide insight into the actions undertaken over time by major manufacturers and retailers. As was called for in the CCHM study, it is time to investigate dynamic, multiperiod strategies.

Second, major sources of asymmetries are missing from the CCHM equilibrium analysis. There are two main sources of asymmetries. They can arise from stable, cross-competitive effects, but can also arise from temporary differences in marketing offerings. One brand on sale by itself might gain much more than if it were promoted along with four other brands in the category. While the CCHM study incorporated measures of distinctiveness into their development of methods for reflecting asymmetric competition, the equilibrium analysis used a simpler model that did not account for this source of asymmetries.

Third, the CCHM effort studied market share, while we here 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 requirements of most economic models. One major alternative to mathematical exploration is multi-period simulations, such as Axelrod's first tournament [Axelrod 1984] or the Fader/Hauser tournament [Fader & Hauser 1988]. While these have the advantage of allowing strategies to be played out over time, they have previously only been undertaken with symmetric and hypothetical market-response functions. We want to use asymmetric market-response functions that characterize brand behaviour in actual markets to study the evolution of robust strategies.

Data from an asymmetric model of a regional U.S. coffee market are used to breed simple artificial agents. *We shall demonstrate that, in the limited tests we can feasibly conduct, these agents outperform the historical actions of brand managers in this regional market.*

2 Modelling the Managers

Competitive marketing strategies can be represented as sets of rules that map states of the market to actions undertaken by brands, brand managers or retailers. These *sets of rules*, in turn, can be represented as chromosome-like *strings*. The *fitness* of each string can be judged by the profits it produces over a period of many interactions, following Axelrod [1987].

A player choosing a strategy can be thought of as choosing a *machine* (a *finite automaton*) or *artificial agent* that will play instead of the player [Marks 1992a]. Such a machine is designed to have a unique action in response to each possible state.[1] The *state* is defined by the history of actions taken by the player and the historic actions and reactions of other players. This line of reasoning builds on developments by Axelrod and Forrest [Axelrod 1987]. They view *players* (e.g., managers) as being characterized by *bounded rationality* [Simon 1972], in which *memory*, *computing ability*, or *competence at pattern recognition* is limited. The *states of the market* are the number of past actions of all players in *limited*

[1] This is a pure-strategy machine (i.e., a strategy chosen with probability 1.0); no mixed strategies are allowed.

memory. If there are p players, a possible actions per round, and m rounds of memory, then the number of states is a^{mp}.

The Axelrod and Forrest study demonstrated that genetic algorithms (GAs) could take the place of the human programmers used in the original Axelrod tournament [Axelrod 1984] or the Fader & Hauser [1988] tournaments. Axelrod reports that the GA evolved strategy populations whose median member resembled *Tit for Tat* and was just as successful. In some cases the GA, which does not require well-behaved, differentiable, globally-convex objective function, was able to generate highly specialized adaptations to a specific population of strategies for particular situations that performed substantially better than *Tit for Tat*.

After Axelrod's pioneering study, other applications of GAs to economics have appeared [Miller 1989; Eaton & Slade 1989; Marks 1992a, 1992b; Marimon, McGrattan & Sargent 1990; Arthur 1990; and Arifovic 1994], with one application in marketing [Hurley, Moutinho, & Stephens 1994].

Our challenges are (i) to develop strings that represent real strategies in asymmetric markets, and (ii) to calibrate asymmetric market-response functions that translate the market states into fitness measures for each brand.

We can coevolve artificial agents, using the asymmetrical profit functions, and then take each of the coevolved agents in the final generation and separately play it against the actual history of the other $n-1$ brands, and assess its performance against that actually achieved by human brand managers. That is, we can ask if our procedure of encoding, breeding and testing has evolved a strategy for Folgers (say) which would have been more profitable than Folgers was historically at competing in the retail coffee market.

3 Asymmetric Competition in a Regional U.S. Coffee Market

3.1 Choice of Market Example

We want to work with an example of competition that exhibits four aspects of real-world markets:

(I) *Differential effectiveness of marketing-mix instruments across brands.* Each brand may have its own unique sensitivity to consumer response to its marketing actions.

(II) *Stable cross-competitive effects.* Some brands gain much more from the losses of certain rivals than would be dictated by market share alone, while other brands are far more insulated by competitive boundaries than the symmetric-market hypothesis would allow.

(III) *Asymmetries due to the temporal distinctiveness of marketing actions.* That is, representing the role of choice context on what brands are chosen: marketing actions must be distinctive to be effective.

(IV). *The dramatic swings in volume that characterize promotion response.* Scanner data reveal that, when viewed at the store or chain level, market response to tactical market-mix decisions is abrupt and dramatic.

The retail coffee market analyzed in Cooper & Nakanishi [1988] satisfies all four criteria. There are eight brands: Folgers, Maxwell House Regular, MH Master Blend, Hills Bros., Chock Full O'Nuts, Chase & Sanborne, Yuban, and an aggregate of premium brands called All Other Branded (AOB). The data track the sales impact of price per pound (net of coupons redeemed), major newspaper ads, in-store displays, and store (not manufacturer) coupons, for 80 weeks in three grocery chains operating in this two-city market. For the sake of simplicity, however, we focus on the 52 weeks of data for Chain One.

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 & Nakanishi 1988, pp. 219-257]. In order to use this simulator as an instructional device, manufacturers' unit costs and promotional costs have been estimated for each brand. This allows us to estimate total profits for each brand for any market scenario. These estimates are thought to be roughly accurate.[2]

Typical behaviour of some brands is to cut their price and engage in newspaper advertising, in-store displays, and coupon distribution after a period of higher prices and no other activity. 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—this is a strategic interaction. The overall patterns of prices and sales for the three major brands available in Chain One (Folgers (F), Maxwell House Regular (M), and Chock Full O'Nuts (C)) are depicted in Figure 1 and the average prices and annual market shares for all brands are shown in Table 1.

There are at least three main ways we might breed artificial agents.

(I). Breed populations of each of the eight brands against the history of the other seven for *each* of the 52 weeks. The procedure would be repeated for each of the eight brands. While this procedure will quickly breed agents to maximize profits against the fixed moves of the other seven in any week, it is essentially static and ignores the multi-period nature of strategic interactions.

(II). Breed populations of each of the eight brands against the history of the other seven over the time frame, with g agents each playing against the entire 52-week period, until convergence. This approach is better, since each brand's g agents are exposed to 52 weeks of the other brands' actions. But the 52-week pattern is still static in that the focal brand's competitors do not react to the actions of its artificial agent, they simply repeat history.

---

[2] 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 these estimates are inaccurate, the validity of our results for the coffee market may be reduced.

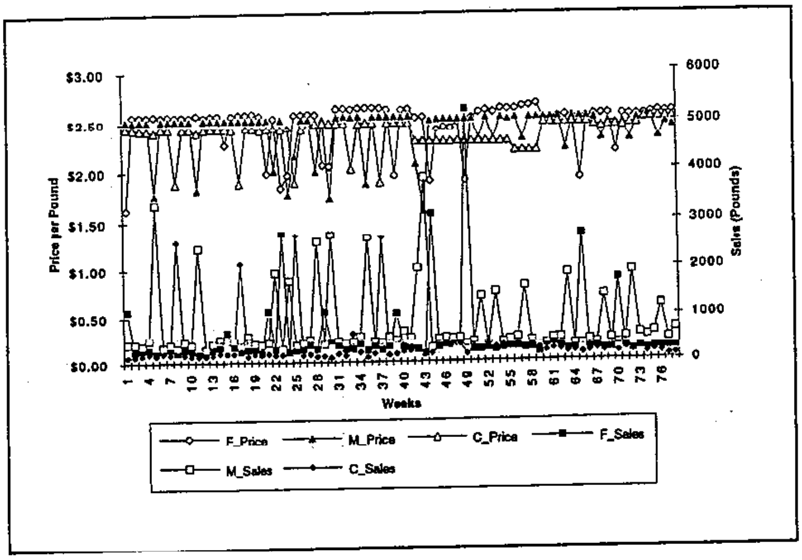

Figure 1. Prices and sales of the three strategic brands

Table 1. Average prices and annual market shares

Brand	Price per pound	Market share
Maxwell House Regular	$2.40	37%
Folgers	$2.45	24%
Chock Full O'Nuts	$2.36	16%
MH Master Blend	$2.78	13%
Chase & Sanborne	$2.36	2%
Hills Bros.	$1.91	2%
Yuban	$3.13	<1%
All Other Branded	$2.56	6%

Furthermore, there is no way around the static nature of the data, since they do not reveal what the contingent strategies of the competing brands might have been. As these contingent strategies are what we are trying to evolve, we believe breeding our agents against historical actions is not adequate.

(III). Coevolve populations of each of the eight brands against all of 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. This is analogous to breeding the agents in a laboratory experiment rather than the field, as in (I) and (II) above. We would then trial the best artificially bred agents for each brand against the historical actions of the other seven over 52 weeks. This approach reveals the best-adapted brand strategy by comparing the brands' scores against actual profits over the historical periods.

Two tests of the artificial agents are explicit in the third approach. One is their profit performance against other artificial agents in the laboratory, the other is the field test of each against the historical actions of the others. Neither of these is perfect, the laboratory test because it is entirely artificial and moreover because convergence of behaviour and genetic drift result in a smaller number of states and so a smaller number of positions on each string being tested for, the field test because it suffers from the lack of learning noted in (II) above. But the only better tests we can currently envisage are to play an artificial agent against the future actions of coffee brand managers either in a brand management game or in the real market. We have not yet conducted such tests.

There are significant problems of complexity with an eight-brand example, especially if a wide range of possible actions are allowed, and hence a large number of possible states of the game need to be encoded for in an agent's bit string. 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 (with a maximum of 51 distinct states) and therefore their fitness never tested. For these reasons we sought to simplify the problem.

3.2 Modelling the Coffee Players

We want to reduce the number of possible states for computational reasons, and, more importantly, for data reasons. We can do this by reducing the number of rounds of memory, which is probably not realistic, by reducing the number of actions of the players (again, not realistic), and by reducing the number of strategic players (again, not realistic). 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?

First, we assume that the decision to use coupons is simply a decision to lower price (which is net of coupons). Rather than considering price to be a continuous variable, with a consequently very high number of states, we represent four price levels. Figure 2 shows that the smoothed frequency polygon for Folgers' prices has four rough peaks. The right-hand or most common peak relates to the shelf price of Folgers, while the others denote promotional prices. The frequency polygons for other major brands have similar quadrimodal characteristics.

Given that each brand has a choice of four prices and also whether to *display* or not, and to *feature* or not, there are 16 possible actions per week per state. In the historical data, features and displays only occur with low prices, and therefore we might reduce the number of actions per brand per week to four, where each price level had an associated feature and display value. Four actions can be coded in two bits, considerably reducing the complexity of the problem.

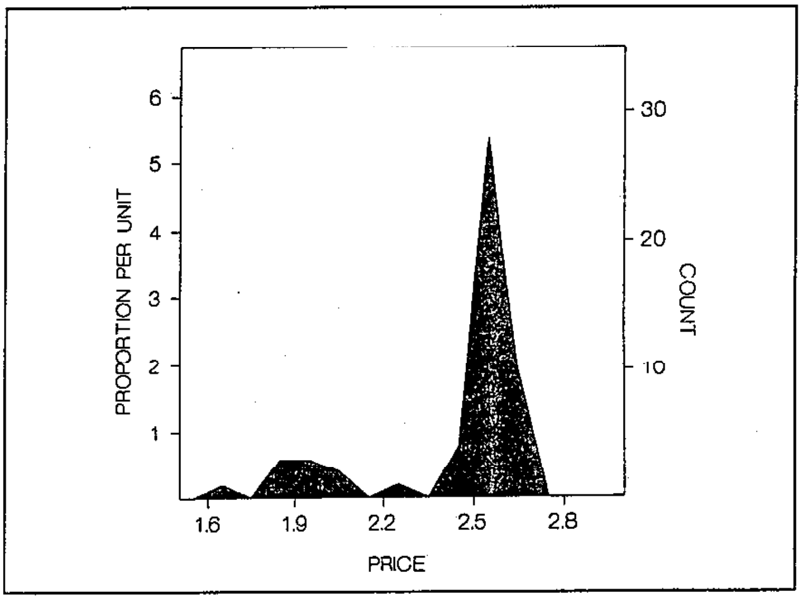

Figure 2. Folgers' price distribution

We model the market as having three strategic players (Folgers, Regular Maxwell House, and Chock Full O'Nuts), with the other brands as fringe players, who act as non-strategic price takers. This means there are only 64 possible states (three players, each with four possible actions) which 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 three strategic brands emphasized in this simplification are by far the major players in this market.

We used a version of the GA[3] to simulate the actual behaviour of the brands in a realistic manner. To reduce complexity we set up the algorithm using a single

[3] We adapted GAucsd, the U.C. San Diego version of GENESIS, originally written by John Grefenstete at the U.S. Naval Research Laboratory.

population of strings for the three brands rather than three separate populations. With coevolution, we did not use the historical pattern of actions, but only the payoffs (profits) as estimated by the *Casper* model, which were used to derive a 4∞4∞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 co-operative or collusive price, a *high* price to approximate the two-person coalition price, a *low* price to approximate the non-cooperative, Nash-Cournot price, and a *Low* price to approximate the envious price. We had also to determine the amounts of feature and display promotions associated with each price level. (See Midgley Marks & Lee [1994] for details). See Table 2 for the marketing mix associated with each action for each brand. The non-strategic sellers' prices per pound are: MH Master Blend $2.90, Hills Bros. $2.49, Yuban $3.39, Chase & Sanborne $2.39, and All other brands $3.68.

Table 2. Possible actions for each strategic brand

Action	**Price** (**$/lb**)	**Feature** (**% stores**)	**Display** (**% stores**)
Folgers			
Low	$1.87	79	68
low	$2.07	82	53
high	$2.38	0	0
High	$2.59	0	0
Maxwell House Regular			
Low	$1.96	95	68
low	$2.33	84	0
high	$2.46	0	0
High	$2.53	0	0
Chock Full O'Nuts			
Low	$1.89	100	77
low	$2.02	99	64
high	$2.29	0	0
High	$2.45	0	0

Each brand participates in 50-round games, with all possible combinations of the other two brands. Although the number of rounds is fixed, the one-round memory eliminates end-game strategies. With a population of size 25, testing each generation of strings requires 8,125 50-round games (325 games per string per generation). Each brand has complete information on all previous actions, but not on other brands' profits (payoffs).

4 Results

4.1 First Experiments—Unconstrained

The first computer experiments found convergence, with all brands pricing at their *Low* price with promotions—not 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 (if not all brands are at *Low* prices), rather than through increased consumption. At least for the period we have data for, we can consider coffee as a mature category with stable long-term consumption rates.

4.2 Second Experiments—Institutional Constraints

To increase realism, we added some institutional constraints. Chain 1 does an excellent job, long run, of maximising 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 player could follow one week's *Low* with another *Low*, and only one player per week at *Low*. Ties of two or more strings (brands) that, 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*.

These constraints resulted in an interesting pattern of behaviour in which brands roughly alternated in pricing *Low*, with the other two brands pricing *low*, *high*, or *High*. But too frequent pricing of *Low* and *low* results in saturation of demand.

4.3 Third Experiments—Demand Saturation

To make the experiments even more realistic, we introduce 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. For details of the demand-saturation implementation, see [Midgley Marks & Lee 1994]. Two things follow if the degree of saturation is greater than 100%: the total sales volume for the latest week is reduced by the degree of saturation, and the profits of the brands are reduced for each of the three competing brands.

With institutional and demand constraints in place, two patterns of competition evolved. In some cases we got convergence to all *low* pricing. In other cases we got patterns of behaviour similar to that observed historically in Chain 1. Figure 3 shows the simulated behaviour of the three strategic brands with the institutional and demand constraints.

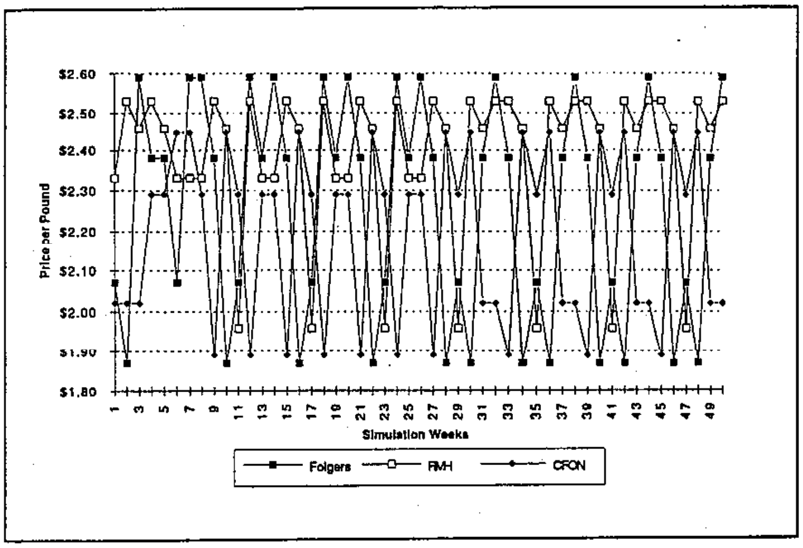

Figure 3. Price paths for the three artificial competitors

It is important to note that the results shown in Figure 3 are for three optimized (coevolved) agents competing against each other over fifty 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 thirty-seven weeks out of fifty, whereas the brand managers for Folgers only promoted fourteen weeks out of fifty. Similar statistics for Maxwell House are thirty weeks out of fifty for the artificial agent, versus eleven weeks in the data. For Chock Full O'Nuts the artificial agent promotes thirty-seven weeks out of fifty, versus seventeen weeks in the data. Over the three brands the artificial agents reduce their prices approximately 2.5 times as often as we observe in the market. If we focus on deep price reductions, the artificial agents employ these 1.9 times as often as we observe in the market. In itself this 'over-competition' is not unexpected, as our artificial agents do not face the practical barriers encountered by brand managers. In the artificial laboratory, information on competitors' actions is received instantaneously, and promotional responses can be implemented within one or two weeks.

In the course of the 'laboratory' simultations the best performing string improved over the 325 50-round games per string per generation by 1.4 times. The best string emerged in the 63rd generation, and remained unbeaten (in terms of its average profit) for the next 37 generations.

4.4 Fourth Experiments—Tests Against Historical Actions

The final series of experiments is not concerned with breeding better agents as such; rather, we took the best agents from the third series of experiments and tested each in turn against the historical actions of their two major competitors[4]. We did this by taking an artificial agent, assigning it to one of the three brands, and allowing it to respond to the historical actions of all other brands over a 52-week period[5]. In fact, as the GA was set to evolve a population of 25, we had 25 'best' agents, and so the test could be repeated 25 times. How well do these strings perform in comparison to human brand managers?

Figure 4 shows that when the final generation of agents is assigned to the Folgers brand most of them do markedly better than human brand managers (as measured by Folger's historical average profit over the 52 weeks). Indeed we have also placed a control line of 25% better than history on the figure and it can be seen that 14 of the 25 agents exceed this. Even the two worst agents generate average profits of 96% and 93% of the historical figure, whereas the best agent does over 240% better than the human brand managers. Although not detailed here, similar results can be generated for Chock Full O'Nuts and Maxwell House, whose best agents do 233% and 120% better than human brand managers do.

While the historical test is limited, in that the competitors do not learn from the changed actions of the brand managed by the artificial agent, these results are impressive. They demonstrate that the 'laboratory' results can be translated to the field. Moreover, given the simplicity of the agents (one-round memory, limited to 4 actions) it is remarkable that they can out-perform human managers. Before we discuss the reasons for this performance, we should ask what the patterns of agent behaviour are that lead to improved profits. This is not an easy question to answer because of the difficulties of presenting all the data in an understandable form. But Figures 5 and 6 shed some light on the issue. Figure 5 shows the historical price actions of Folgers compared with the price actions of the best agent (string 24). Figure 6 shows the same historical actions compared with the worst agent (string 20 of the final generation).

---

[4] With the historical actions of the other five brands input to the profit calculations but not 'recognised' by the agent: the perceived market state is invariant to the other five's actions.

[5] 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, using Figure 2 for Folgers, etc.

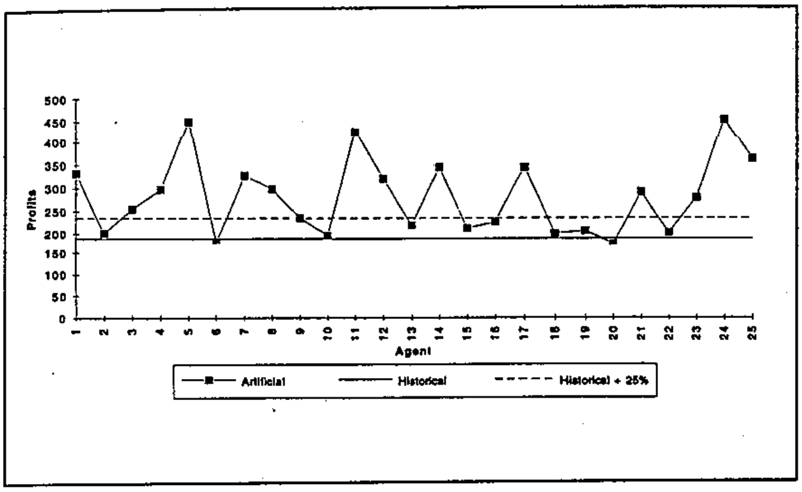

Figure 4. Profits for Folgers: artificial agents versus history

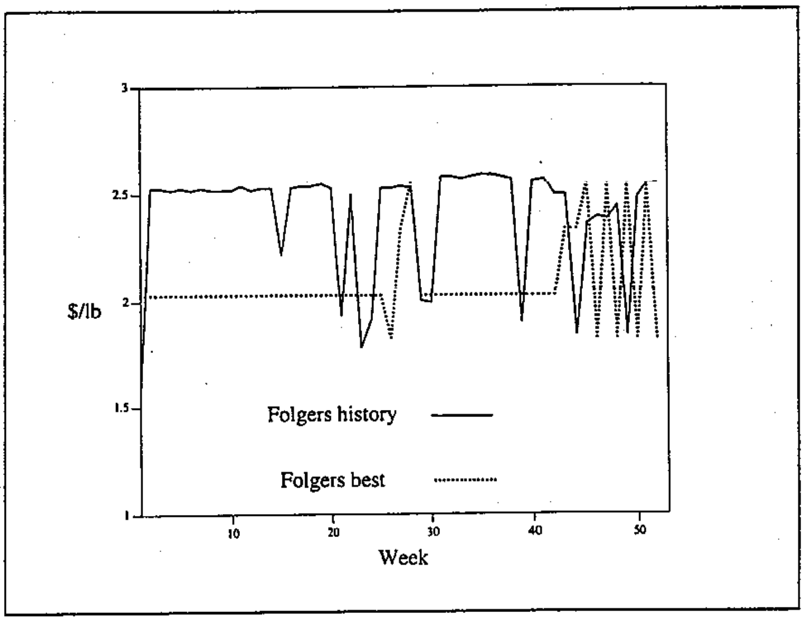

Figure 5. Folgers' price paths—best agent versus history

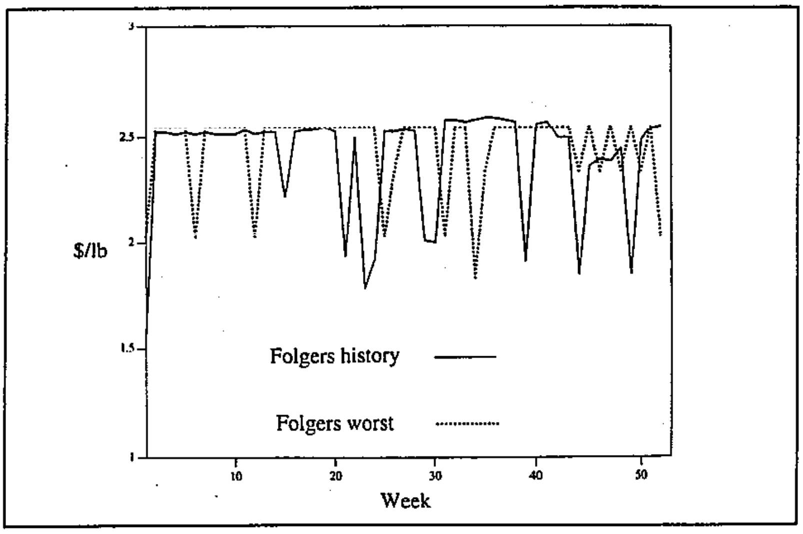

Figure 6. Folgers' price paths—worst agent versus history

The comparison between these two figures suggests that while the "worst" agent behaves quite 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.

5 Conclusion

The general conclusion is that the artificial agents price promote more frequently than human managers. We observe the highest level of promotion when the three optimized agents are competing with each other in our 'laboratory'—an environment which perhaps represents the most competitive scenario we can achieve. Indeed, we might define these results as the maximum competitive intensity possible in this market (given our sales model and institutional constraints). All actual markets would be likely to show less intense competitive activity. Hence, it is not surprising that when we place one of these optimized agents back into the historical market we observe a lower frequency of promotion. This is the case for many of the final strings—whose behaviour more resembles human managers. But it is still true that the best of our agents promote more frequently than do their human counterpoints and we can speculate on the reasons why this might be so.

One reason may be that human brand managers are not in a position to respond to competition on a week-by-week basis. More likely, they negotiate with the chains for a series of promotions to occur across a defined promotional period (often of thirteen weeks duration). Major responses to competitive actions then occur in the next promotional period, rather than by immediate adjustments to the current promotional plan. This suggests that competitive response in real markets may be more measured and less immediately reactive than that generated by our optimized artificial agents. *Institutional constraints may therefore serve to dampen competition.*

But there may be reasons for the greater level of promotion which have more to do with our agents than with the institutional constraints and brand managers: the choice of one-round memory and the selection of the four reference prices. One-round memory restricts the agent to only being directly 'aware' of the most recent actions of its competitors. Two or 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 market-place behaviour. For example, observing that a competitor has promoted for two consecutive periods implies greater aggression than if that competitor has only promoted for one period out of two.

What then are the managerial implications of this approach? We believe these are threefold. First, the artificial agents allow the managers of any brand to check future 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, the agents also enable managers to test 'what-if' scenarios, both for their own brands and for the brands of their competitors. Both these may help alleviate the resistance to market modelling which is observed in many consumer product companies. In our opinion some part of this resistance stems from the static or competitively myopic nature of current modelling approaches. Managers expect models to be able to simulate the consequences of a planning period (often four promotional periods or a year) and to factor in likely competitive responses. Third, the agents may be useful in training junior brand managers: the agents could form the basis of a game whereby junior managers make decisions for one brand and the agents for other brands provide the competitive test of these decisions. With appropriate agents this would inject an element of realism into training by simulation games. This element is missing from many games at present because they use other teams of junior managers to make the competitors' decisions and also often have unrealistic algorithms for market response.

Acknowledgements

The authors wish to thank 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. This research is supported, in part, by the Anderson Graduate School of Management, 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 TIMS Marketing Science Conference in St. Louis, Missouri, the Society of Economic Dynamics and Control Conference in Nafplion, Greece, and the Australasian meeting of the Econometrics Society in Armidale, N.S.W.

References

Arifovic, J. (1994): Genetic algorithm learning and the cobweb model, *Journal of Economic Dynamics and Control*, 18: 3–28.

Arthur W.B. (1990): A learning algorithm that mimics human learning, Working Paper 90-026, SFI Economics Research Program, Santa Fe Institute.

Axelrod R. (1984): *The Evolution of Cooperation*, Basic Books, New York.

Axelrod R. (1987): The evolution of strategies in the iterated Prisoner's Dilemma, *Genetic Algorithms & Simulated Annealing*, L. Davis (ed.), Pittman, London.

Carpenter G.S., Cooper L.G., Hanssens D.M., and Midgley D.F. (1988): Modeling asymmetric competition, *Marketing Science*, 7(4): 393–412.

Cooper L.G. and Nakamishi M. (1988): *Market-Share Analysis: Evaluating Competitive Marketing Effectiveness*, Kluwer, Boston.

Eaton B.C. and Slade M.E. (1989): Evolutionary equilibrium in market supergames, mimeo., November.

Fader P.S.and Hauser J.R. (1988): Implicit coalitions in a generalised Prisoner's Dilemma, *Journal of Conflict Resolution*, 32: 553–582.

Hurley S. Moutinho L., and Stephens N.M. (1994): Applying genetic algorithms to problems in marketing, *Proceedings of the 23rd EMAC Conference, Marketing: Its Dynamics and Challenges*, ed. by J. Bloemer, J. Lemmink, H. Kasper, European Marketing Academy, Maastricht.

Marimon R., McGrattan E, and Sargent T.J. (1990): Money as a medium of exchange in an economy with artificially intelligent agents, *Journal of Economic Dynamics and Control*, 14: 329–373.

Marks R.E. (1992*a*): Repeated games and finite automata, *Recent Advances in Game Theory*, ed. by John Creedy, J. Eichberger, and J. Borland, Edward Arnold, London, pp.43–64.

Marks R.E. (1992*b*): Breeding optimal strategies: optimal behaviour for oligopolists, *Journal of Evolutionary Economics*, 2: 17–38.

Midgley D.F, Marks R.E, and Lee L.G. (1994): Breeding competitive strategies, UCLA Anderson Graduate School of Management Working Paper No. 242.

Miller J.H. (1989): The coevolution of automata in the repeated Prisoner's Dilemma, Working Paper 89–003, SFI Economics Research Program, Santa Fe Institute, July.

Simon H. A. (1972): Theories of bounded rationality, from: *Decision and Organisation: A Volume in Honor of Jacob Marschak*, C. B. McGuire & R. Radner (eds.), Chapters 8 & 9, North-Holland, Amsterdam, pp.161–188.