

**USING EVOLUTIONARY COMPUTING TO EXPLORE COMPLEX ECONOMIC  
PHENOMENA: MODELING THE INTERACTIONS BETWEEN CONSUMERS, RETAILERS  
AND BRAND MANAGERS**

**By**

**Robert E. Marks**

**Australian Graduate School of Management**

**University of New South Wales**

**Australia**

**bobm@agsm.edu.au**

**&**

**David F. Midgley**

**INSEAD**

**France**

**david.midgley@insead.edu**

---

The authors acknowledge the financial support of the AGSM and INSEAD and would like to thank *Information Resources, Incorporated* for the data and Peter McBurney, University of Liverpool for helpful comments on an earlier version of this paper.

## ABSTRACT

Competition among brands that sell through a supermarket chain is a repeated moderated oligopoly, where the retailer moderates brand competition through its choice of which sales promotion to allow at its stores in any 7-day period. It does so to maximize its profits by increasing its sales volumes not only in the focal category but possibly also in other categories. This moderated competition is complex, and we use an evolutionary algorithm — the Genetic Algorithm — to model rivalry among the brands and to examine three models of retailer behavior— zero-intelligence (ZI), naïve, and sophisticated choice of sales promotions. Our goal is to establish whether there are positive returns to more sophisticated moderation. We find that, although the brand behavior is more realistic with the sophisticated retailer, on average the category volumes are lower than with either of the naïve or ZI retailer. The dominant brand, however, does benefit with the sophisticated retailer, by preserving its profits and reducing those of its competitors. We speculate there may be side-payments or sufficient spin-off benefits in other categories to compensate the retailer for the loss of volume in the focal category.

## I. INTRODUCTION

The profit-maximizing strategies of economic agents operating in perfectly competitive markets — in which there are many buyers and sellers, each of which is a price taker — are well known. So are the strategies for pure monopoly (or monopsony) — in which there is a single seller (or buyer). In contrast, the strategies for agents operating in oligopolistic markets, with few sellers offering close substitutes for sale to many buyers, are less clear. There are many Nash equilibria in the strategic rivalry among the sellers in an oligopoly, the simplest of which are the Cournot (quantity) and the Bertrand (price). Yet, oligopolies are more common in reality than monopolies or perfect competition, and appropriate strategies for oligopolists therefore of greater interest to both scholars and managers.

Our paper continues a program of research into such oligopolies. In particular, we study the repeated interactions between small numbers of sellers in a retail market for ground coffee. In published work we have previously sought strategies for each brand that improve the brand's weekly profits compared with historical results for an actual market. The brand managers use several marketing instruments to encourage sales: price, coupons, aisle displays, and advertising in the retailer's newspaper. An important characteristic of these real-world markets is *asymmetry* (Carpenter, Cooper, Hanssens and Midgley 1988). Brands have different costs of production, they hold different positions in the market, and consumers react differently to brands' marketing instruments. This occurs for many reasons, including history, patterns of competitive reaction, and consumer perceptions. Our approach has been to use econometric models of consumer response to marketing instruments, coupled with bit-string agents modeling the decisions of brand managers. We model adaptation of brand-agents by use of an evolutionary technique, the Genetic Algorithm (GA). By analyzing the behavior of these agents, and by contrasting it with historical brand behavior, we use the GA to explore agent learning in a more realistic setting than is available for closed-form solutions to oligopoly behavior. An eventual goal of our research program is to help brand managers and retailers manage their interactions more efficiently.

In previous papers (Midgley, Marks and Cooper 1997, Marks, Midgley and Cooper 2006), however, we largely ignored the role of the retailer — here, a supermarket chain with many branches — in

the interaction between brands and consumers. . The only role we allowed the retailer previously was to enforce a policy commonly seen in such supermarket chains, namely that only one brand can run a price promotion in a category in any week and only for one week.

In sharp contrast, in this paper we allow the retailer to be an *active moderator* of the rivalry among the brands, as it attempts to maximize its profit too. Now, it could rightly be argued the retailer's focus is on the total profit it makes from the category as a whole, not just from one or more brands. But, because of the asymmetry mentioned above, this total category profit turns out to be dependent on which brand promotes when and to what extent, as we show below. Furthermore, in some categories *advertised* promotion of a prominent brand can attract shoppers to the store and, because people usually buy in several categories when they shop, have broader impact on that store's sales. Seen from that perspective, store promotions are also a means of competing with rival chains. Overall, the retailer therefore has good reason to exercise care in choosing which promotions run in any week. In ignoring this choice, our previous models had an important limitation.

In this paper we examine three levels of moderation by the retailer: (1) zero-intelligence (ZI) decision-making, (2) naïve decision-making and (3) sophisticated decision-making, as discussed below.

#### *A. Zero-Intelligence Retailers*

Our ZI (stochastic-choice) retailers are not new in concept, although they have not been applied to a marketing model before. In finance, ZI traders were first assumed by Gode and Sunder (1993) and have also been used in computer science (e.g. Cliff 2001). An extensive discussion of ZI traders can be found in Duffy (2006). An important finding is that in some settings these ZI traders outperform human traders. Here we use them as a base case to compare with other levels of moderation. If two or more brands offer to promote in the same week, the supermarket chain's choice is made by a simple coin toss, or equivalent.

#### *B. Naïve Retailers*

Our naïve retailers are more novel. They represent a supermarket chain which chooses between brand promotions simply from the price discount offered in any week, the type of promotion and the proportion of stores in the chain that will run the promotion. These are the explicit features of the 'deal' offered by

the brand manager. Naïve retailers ignore asymmetries between brands: they assume an average impact; ignoring the fact that one brand may have twice the promotional impact of another. In defense of what may seem extreme naivety, please note that supermarket decision-makers deal with thousands of product categories, each category with several brands and sizes, all of which may differ in promotional impact. How much knowledge can they have of any one category?

### *C. Sophisticated Retailer*

Finally, and in contrast to the ZI and naïve retailers, our sophisticated chains thoroughly understand the coffee category, including asymmetries across brands and types of promotion and consumer stockpiling. Our sophisticated retailers also have some knowledge of events in competing supermarket chains. All in all, this represents a retailer whose analysts have thoroughly studied the coffee category using time series data and econometric techniques. Such retailers exist, although they are not perhaps typical.

### *D. Our Hypotheses and Preview of Results*

In our discussion of retail moderation, there is a hypothesized hierarchy of increasing performance, from ZI to naïve to sophisticate. Despite these expectations, we find that ZI retailers perform well. They did in the financial trading case too, but nonetheless we were surprised, given our different setting. The other striking results are threefold. (1) Sophisticated retail moderation lowers category volume and brand profits (except for the leading brand); (2) but results in much more realistic behavior for our brand-agents compared with real managers; and (3) because of brand asymmetries, these brand-agents evolve to avoid head-to-head competition for the retailer's choice.

### *E. Structure of the Paper*

In the next section we briefly outline our prior work, before showing how we develop and implement the naïve and sophisticated retailer models. We then outline the literature on ZI agents and formulate our hypotheses. Then, in the methodology section, we present our Monte Carlo approach, the range of simulations that we ran, and show how we analyze the results of these simulations. Finally, we present the results of these analyses, discuss these in more depth, and close with our conclusions and suggestions for further work.

## II. PRIOR WORK

Our broad interest is in the effects of asymmetric market response on the competitive actions of managers; in particular, how managers will compete in multi-period games where each brand elicits a different response from consumers and each firm faces different costs. This is the world faced by managers of the product categories sold in supermarkets. We understand consumer response in this space, and we have methods for modeling these responses (Cooper and Nakanishi 1993). We know much less about the competitive actions of managers, and there have been few attempts to model the repeated interaction evident in these settings, mainly because analytical solutions to these interactions are intractable.

In Midgley, Marks and Cooper (1997) we showed how we can model the actions of each brand manager as an outcome of a stimulus-response finite automaton playing a repeated game. Using the Axelrod and Forrest (1987) representation of these artificially adaptive agents as bit strings, we employ a GA and an existing consumer-response model—Casper (Cooper and Nakanishi 1993)—to co-evolve artificial agents for managing brands of coffee in a regional U.S. market. With the Casper model we can predict the weekly sales, revenues and profits for each brand, given any combination of price, store display and advertising for all the brands on the market. In other words, if we know the actions of brand-agents, then we can predict their profits. Casper performs well, achieving a squared cross-validity correlation of 85% in a hold-out sample (Cooper and Nakanishi 1993).

While this response modeling is well-known, our introduction of agents for the brand managers was novel. To do this we first analyzed the historical behavior of brand managers to discover a few typical marketing actions. An ‘action’ is a specific combination of price, advertising and display. For example, a major promotion for a brand would be a discounted price of \$1.89 together with in-store aisle displays and advertising in the store’s newspaper. A non-promotional action would simply be the normal or “shelf” price, for example \$2.55. This shelf price can also vary from week to week without necessarily being promoted or brought to the customers’ attention before they buy. In reality, brand managers use many such combinations, although there are some typical patterns.

To make the problem more manageable, we use cluster analytical techniques to define a smaller set of typical actions. We outline these techniques later. We can then give the bit-string agent representing each brand a small menu of actions from which to choose. And we can partition the previous actions of competitors into a small number of equivalent states. The job of each brand-agent is then to select an action from this menu that will be profitable for it in the next period of the game, given the stimulus of its perception of the previous state of the market. In our previous work, we used simple agents limited to four actions, with perceptions restricted to four equivalent states for a single previous period of the game (Midgley, Marks and Cooper 1997, Marks, Midgley and Cooper 2006). This resulted in agents whose mappings (encoded selections) from possible states to best actions have a 134-bit representation. The number 134 comes from modeling three brands each with four possible actions ( $2 \times 4^3 = 128$ ); we also use 6 bits of ‘phantom memory’ to endogenize initial conditions (Midgley, Marks and Cooper 1997). The GA then ‘evolves’ mappings that maximize profits over many multi-period games, pitting various agents against their competitors. The GA achieves this by *selecting* and *hybridizing* the better performing strategies from each generation of these games to create more profitable strategies for the next generation, an iterative cycle that continues until no further improvement is possible.

We found that these hybridized agents performed well both in these interactions and when we pit a single agent against the historical actions of human brand managers. The latter test is ‘unfair’—there is no opportunity for the human to respond to the agent— but this was nonetheless a surprising result, given the simple agents we used. Moreover, in developing these agents, we learned that store policies (in particular one promotion each week) and demand saturation were important to incorporate in our model (Midgley, Marks and Cooper 1997, Marks, Midgley and Cooper 2006).

### **III. THE RETAILER AS AN ACTIVE MODERATOR OF PROMOTIONS**

#### *A. Limitations of Our Prior Work*

In our prior work the retailer imposes the policy of one promotion each week but makes no choice of which promotion this should be. If only one brand agent offers a promotion, it is accepted. That is reasonable: any promotion adds excitement and expands category volume to at least some extent. Further,

the retailer receives discounts and incentives from the brand, which protect against losses. The more interesting case is if two or more brands offer a promotion for the same period. In our prior work the retailer simply tosses a coin to decide which brand will promote. In doing so the retailer takes no account of the features of the competing promotions. This is clearly a major limitation. The retailer is missing an opportunity to choose which promotion will maximize its profit or volume. To overcome this limitation, this paper extends the retailer decision-making to the naïve and sophisticated levels mentioned before, while keeping the previous coin-tossing or ZI level as the base case.

### *B. Developing Naïve and Sophisticated Levels of Retail Moderation*

We have 78 weeks of data from our focal supermarket chain, with similar data from competing chains. They were made available to academic researchers in marketing by *Information Resources, Incorporated*. These data include sales, prices and marketing instruments for the eight brands of coffee regularly sold in all the chain's stores. Previously we chose to focus on the three main brands, namely *Folgers*, *Chock-Full-of-Nuts (CFON)* and *Maxwell House* (Midgley, Marks and Cooper 1997). Here, we have added a fourth brand, *Hills Brothers*, an independent niche player that represents a useful contrast to the main brands. These four brands account for roughly 75% of the volume sold in the focal supermarket chain, and *all* the promotional activity. The remaining four brands are passive players that never promote and do not change prices often. For that reason we do not include them in the retailer models, although they are passively present in the consumer response model. Only the four “active” brands are included in the retailer models.

The dependent variable we choose for these models is *total category volume* (pounds per week). Experts in the field suggested this variable both as a reasonable proxy for retailer profit (which we do not have) and for the impact on sales of other categories in the store. Again, please note that for promotions the retailer gets offsetting discounts from the brand managers on the wholesale price, as well as specific payments for each store for the costs of displays and newspaper advertising. While the price to the consumer is reduced, the retailer's profit margin is probably unchanged. Total category volume is thus a

useful criterion for deciding between two competing promotions. The retailer should prefer the promotion with the greatest impact on weekly sales.

The independent variables in the data include promotional and shelf prices for every brand. They also include the proportion of stores in this chain displaying or advertising any promotion (expressed as a proportion of chain volume, to weight suitably for different store sizes). From these variables we build a predictive model of category volume to use in making the retailer's choice between competing promotions.

*1) Naïve Retailers:* We define a naïve retailer as only concerned with the discount for the promotion, the type of promotion and the proportion of the stores in their chain that will sell at this discount. They are blind to brand. To operationalize this idea we ran a simple OLS regression on total category volume, with no transformations or lags on any variables and no attention to the time-series properties of the data. This model has an R-squared of 65% and the coefficients shown in Table 1.

[Table 1 about here]

An econometrician would have doubts about this model. We do not use knowledge of the setting in developing the model, and it has biased coefficients. Our goal, however, was to represent a retailer who only had a rough idea of how category volume might change, given simple features of a specific promotion. In that sense we think the model is a reasonable one. Note, this retailer also does not distinguish between brands. We should make two comments on these results. First, given its simplicity, the R-squared is surprisingly high. 'Naïve' buys you notable predictive power. Second, the fact the coefficients for major and minor promotions are the same needs interpreting with care. As the range of discounts associated with major and minor promotions differ greatly, the impacts of these also differ.

*2) Sophisticated Retailers:* In contrast, we assume the sophisticated retailer is an econometrician who analyzes and understands the impact of promotions thoroughly. We operationalised this with an econometric model where we log-transform the properties of the dependent variable to better fit the assumptions of regression. We also account for the negative autocorrelation in the time series which the Durbin-Watson test suggests. This we do by standard autoregressive procedures. Analysis also reveals we should include a factor for the impact of sales at another chain on sales at our focal supermarket.

Several model formulations were explored, including models where we transform the independent variables to lessen multicollinearity. We did this by weighting the promotional discount by the proportion of stores in which it was available. These models need an autoregressive coefficient, but no lagged dependent variables. We also examined models with separate discount and display variables. These models included a lagged dependent variable and needed complex autoregressive lags. They gave significantly better R-squareds, but have the same implications as the ones above. However, they also have more multicollinearity (because the depths of discount correlates highly with the display proportions) which means the *incremental effect* of promotions is less well calibrated (larger standard errors on coefficient values). As this incremental effect is our main focus, we prefer the models with the weighted discount. We show the final model in Table 2. This has an R-squared of 88% with 66 degrees of freedom.

**[Table 2 about here]**

This model represents very sophisticated retailers. They incorporate information about prior sales and household stockpiling (the lag coefficient), they understand that promotions by different brands have different impacts, and they understand that activities at another chain have implications for their sales. The main brand asymmetry revealed by statistical tests is that the coefficient for the major *Folgers* promotion differs from those for *Maxwell House* ( $p < 0.01$ ) and *CFON* ( $p < 0.05$ ). Overall, the model incorporates a notable amount of asymmetry and has high predictive power. The coefficients shaded in Table 2 are those relevant to a choice *between* promotions. Before we can use this model to make choices, however, we first need to find out what those promotions should be.

*C. Brand Actions*

We extend the number of actions that are available to a brand (and with which to partition competitive actions in previous periods) from the four of our previous work to eight here. This we do in two steps. First, we examine the box plot of the shelf prices of the four focal brands. The median, top and bottom quartile, and top and bottom whisker values of the box plot become the five shelf prices which we allow the brands to use. This covers the range of shelf prices in a consistent and representative manner. Where reasonable, we stylize these prices to end in the digit 9 (which is common in this supermarket chain). The

five shelf prices are \$2.19, \$2.39, \$2.49, \$2.55 and \$2.69. The corresponding actions are labeled as Low2, High1, High2, HIGH1 and HIGH2 in Table 3.

**[Table 3 about here]**

Second, we need to decide the number of promotions and their characteristics. Exploratory analysis reveals there are two types of promotion. One is a major promotion with a significant price discount (using the average price of non-promoted brands as the reference point), store displays and store advertising. The other is a minor promotion, with a smaller discount and only store displays. In the data these occur with a frequency of two major to one minor. As we have three actions remaining of our eight, this suggests two actions to characterize major promotions and one to characterize minor. Therefore, we use cluster analysis to dissect the promotional data, using one more cluster than suggested to understand commonality, heterogeneity and to isolate outliers. For major promotions, this extra cluster turns out to be diverse, essentially a catch-all category, and the other two clusters are much more homogeneous and contain most the observations. So we characterize major promotions with these clusters. They represent two different depths of price discount, with both display and advertising in most of the chain's stores. The values we use for each of the two actions are those from the observation closest to the cluster centroid. These are shown in Table 3 as LOW1 and LOW2. For minor promotions the extra cluster is anomalous in that these promotions occur in few stores. From the data, we believe that they were mistakes! Namely, promotions that some store managers forgot to remove at the start of the following week. So we use the other cluster to characterize minor promotions. This has a moderate discount and is displayed in most stores but is not advertised. In Table 3 it is labeled Low1.

As a final comment here we note that these eight actions are not brand specific. Rather they represent a menu of common actions the agents for each brand can use. An alternative strategy would be to develop brand-specific actions. Indeed, we did so in Midgley, Marks and Cooper (1997). This is, however, both complex and optimized. Common actions are simpler and more defensible. Moreover, because of asymmetries, the same actions made by different brands will result in different outcomes, for the brands and for the retailer, *ceteris paribus*.

#### *D. Choice between Promotions*

From the retailer models and the three promotional actions, using Casper we can estimate the category volumes that result when each of the four brands runs a promotion. By comparing these volumes, our artificial retailer can choose the best one from those offered by the brand-agents. Table 4 shows the predicted volumes that each promotion will achieve (ranked from highest to lowest impact).

**[Table 4 about here]**

Taking the naïve retailer first; any promotion results in significantly increased volume. Indeed, all are around twice the reference non-promotional or ‘normal’ volume. Interestingly, the biggest of the major promotions (LOW1) results in only 9% more sales than does the minor promotion (Low1). But an economically rational retailer would prefer promotions according to their impact, that is,  $LOW1 > LOW2 > Low1$ . For the sophisticated retailer the choice is more complex. There are clearly brand differences, and impact also varies from a promotion that is only 3% more than normal volume (a promotion for *Hills Brothers*) to several that deliver over twice normal volume. The highest impact promotion, a display-only promotion for *Folgers*, is interesting in that it produces the greatest impact, despite being at a lower discount than other promotions for the same brand. And *Folgers* is the only brand where the logical order of promotions breaks: for all other brands they are, in order of impact,  $LOW1 > LOW2 > Low1$ . We speculate that this is to do with *Folgers* being the market leader. Whatever, when two or more brand agents propose that each promotes at the same time, the retailer chooses a single brand to promote in the period according to these category volume rankings. As we have discussed, different kinds of retailer use different methods to choose the promoting brand in any period. But the fate for the brand, or brands, that are not chosen to promote in the period is the same across retailer models: the rejected brands use their High2 stylized actions instead of the desired LOW1, LOW2 or Low1 promotional actions.

#### **IV. HYPOTHESES**

Gode and Sunder (1993) report a program of research into the impact of the institutional structure of markets. They examine the continuous double auctions of the stock market and whether these result in convergence of prices to the equilibrium, market-clearing price. These authors successively assumed

simpler and simpler automated trading agents and compared the results with laboratory results of human subjects. They also compared the efficiency of such market outcomes with the theoretical 100% efficiency possible by exchange between fully informed and rational players. Eventually they came up with the simplest agents, the ZI agents, which randomly choose their stated order prices, with no observation of the state of the market (demands and supplies) and no memory of past states. The simplest of these agents (ZI-U) were unconstrained in their possible order prices. These ZI-U agents can reach an allocative efficiency (sum of trader surpluses over potential) of over 78% (Duffy 2006). A slightly less naïve agent (ZI-C) is budget-constrained in its stated order prices: it may not reduce its book profit. These ZI-C agents can reach an efficiency of over 97%, still choosing order prices randomly (if constrained), with neither perception of current market conditions nor memory of past conditions.

Computer scientists such as Cliff and Bruten (1997) explore not the institutions of the market as Gode and Sunder were doing, but the minimum “intelligence” of agents necessary to achieve equilibrium in a reasonable time and with reasonable efficiency. They used machine learning (GAs) to explore this minimum intelligence in simple continuous double auctions, where each agent preserved a target profit margin that could be adjusted. Their optimized agents they dubbed Zero-Intelligence Plus (or ZIP) agents, which displayed efficiencies and rates of convergence higher than previous studies. Subsequently, others found that ZIP traders could outperform human traders (Das, Hanson, Kephart and Tesauro 2001). Computer science researchers at Hewlett Packard and IBM and elsewhere have continued to study the heuristics of ZI and ZIP agents, but our use of them in marketing is novel.

Here, we model the retailer’s decision of which promotion proposal to choose for the next period. We model ZI, naïve, and sophisticated retailers: the naïfs and the sophisticates try to maximize their expected channel volume, but the ZI just choose randomly.

Rather than comparing our agents with human retailers, we compare the three kinds of retailer against one another. We expect their performance to rise with complexity:

H1: Naïve retailers outperform ZI retailers;

H2: Sophisticated retailers outperform naïve retailers.

## V. METHODOLOGY

### A. Representing brand managers as bit string agents

We model the artificial brand managers as stimulus-response automata. The stimulus is this week's market state (defined by the marketing actions of the four strategic brands). The response is the brand's *proposed* actions for the next week, but the *eventual* actions for each brand are the outcome of the supermarket chain's moderating decision, as it responds to the proposals of the competing brand-agents.

We model each brand-agent as a binary string representing a mapping from all possible states of the market to a response for each of these states. How does this work? With eight possible marketing actions to choose from, we can use a triple of bits on the string to code for the response to each state of the market ( $2^3=8$ ). How many triples are sufficient to represent all possible states? With four brands, each with eight possible actions and one-week memory, the answer is  $8^4 = 4,096$  states, which leads to a string length of 12,288 bits ( $4,096 \times 3$ ). More generally, there are  $a^{mp}$  possible states, where  $a$  = the number of actions,  $m$  = the number of periods remembered and  $p$  = the number of strategic players (Midgley, Marks and Cooper 1997). An extra 12 bits of 'phantom memory' are also needed to encode the agent's initial belief about the current week at the start of the game (a triple for each of the strategic brands). This gives us 12,300 bits per string, a 90-fold increase over our prior work but possible to compute on a current PC such as the Mac G5 Dual 2GHz PPC used for this work. Using strings of such a length has been called a brute-force approach, and may not embody great coding sophistication. But it is straightforward and readily executed on today's computers.

### B. Using the GA to optimize these strings

To make this representation work, we use the GA to explore good mappings from possible states to responses. From a random pattern of bits on each brand's string, the GA searches for patterns that enable the string to perform well in the oligopolistic competition with other brands. This work extends that of Marks (1992), using the ability of the GA to search the disjoint space of strategies, as Fudenberg and Levine (1998) suggested. The GA borrows from our understanding of evolution to search for solutions to

problems not easily solved otherwise (Mitchell 1996). The algorithm is well known. We define a suitable fitness score and create an initial population of solutions. We then use the fitness score to select a subset of better-performing individuals to be the “parents” of the next generation. We also allow the “crossover” of pairs of parents, and flip each bit from zero to one or conversely (“mutate”) with a small probability. The fitness of each member of the new population is calculated and the algorithm repeats until little or no further improvement is obtained. Specific details are available from the authors.

Whereas engineers have used the GA as an optimization tool, social scientists have used it in a different way: as a means of simulating co-evolution. In our model, each brand manager can be thought of as learning from its rivals’ behavior and from its rivals’ responses to its own actions as the generations pass. This mutual leaning means the competitive environment changes, even as each brand-agent learns to compete more effectively. Because of this ‘Red Queen’ effect (Van Valen 1973); there is no necessary increase in fitness, even as the GA winnows the succeeding generations of their worst performing strings.

Co-evolution requires a separate population for each of the strategic players. A single population would allow extra-market communication and learning to occur via the genetic operations of selection and crossover. Not only would this be illegal under antitrust laws, but such social learning (Vriend 2000) is not what we want to model. Necessarily, four separate populations require a much more complex GA program, but only a co-evolving GA is appropriate. We extensively rewrote the GA software, GAucsd (Schraudolf and Grefenstette 1992), to allow the simultaneous simulation of four populations of agents (this code is available on request from the authors). We use a population size of 25 strings, each string being 12,300 bits long, with four populations corresponding to the four brands.

### *C. How is fitness determined?*

We define the fitness score as the average weekly profit of the brand. This definition assumes profit maximization is the goal of the brand-agent and does not allow other strategies that may exist in real markets. For example, cross-subsidization across the various brands owned by one manufacturer, accepting losses to establish a new brand, *et cetera*. So the definition chosen here is more appropriate to

mature markets and brands, and to companies that manage their brands individually. This is the most common marketing situation but it is clearly not the only one.

Each agent-string in its population plays a 50-round game with every possible combination of strings from the other three brands' string populations and its average profit is its fitness. The GA used 25 strings for each of the four populations. Therefore, testing each generation required 73,125 50-round games, or 2,925 games per agent-string per generation. This is a non-trivial simulation, especially as we have found it takes around 100 generations to achieve satisfactory convergence for the four brands, which equates to 29,250,000 50-round games. Need some words here

Why do we determine fitness by playing a string against all other combinations of strings? This is not what real managers can do. With our very long strings, differences in the genotype (the structure of the string) are not of great interest: in parts of the string where the evolutionary pressures are not great there is no reason to expect convergence from the initial random pattern; what we are interested in is differences in the phenotype, the behavior, as we have discussed. – we can think of the strings in a population as thoughts in the head of the manager modeled by the single population: s/he conducts thought experiments (what will happen if I do this, do that, do the other ...) to decide what's best – the fittest string corresponds to the best notion. If we model all managers as doing this, then we need to play each string in a population against all combinations of strings in the other populations. Of course, no manager will be as thorough as this, so our results will probably be unrealistic, but a manager could attempt to do this, and given enough time would get results similar to ours. Moreover, this is how to determine the Nash equilibria – the combinations of best responses, where "best" might be constrained by the bounded rationality of the human brand manager.

Can we seriously claim to search this vast space with a population of 25 and 100 generations? Well, GA researchers have found that there are non-linear interactions among the various parameters of the algorithm: population size, crossover rate, and mutation rate. DeJong (1975) found for his suite of optimization problems that the best population size was 50-100, with a crossover rate of 0.6/pair of parents and a mutation rate of 0.001/bit. For his examples Grefenstette (1986) found that these parameters

should be 30, 0.95 and 0.01 respectively whereas Schaffer et al. (1989) found ranges of 20-30, 0.75-0.95, and 0.005-0.01 respectively. Gathercole and Ross (1997) found that small populations (50) over many generations outperform large populations over few generations. But there is no simple formula for determining these parameter settings; they depend on the problem being optimized. In our work 100 generations appear to be sufficient with a population of 25. This can be seen from the good rate of convergence as the simulation approaches 100 generations.

#### *D. Monte Carlo methods and analysis of outputs*

We also have to consider that a single 100-generation ‘run’ of the GA is not enough, as the outcome may be an artifact of the initial strings we choose to start from. We therefore repeat this whole procedure 50 times, using randomly drawn initial strings, and analyze the outcomes to see the degree of agreement between them. For comparability with our previous work, we focus on the behaviors of the brand agents (the relative percentages of times that they use each of the eight actions) and use cluster analysis to determine agreement between the 50 runs. Other bases for defining ‘agreement’ are possible but not considered here. To overcome problems of differing variance and collinearity, the 32 variables (four players times eight actions) are reduced to a smaller number of canonical variables, using the approximate covariance pre-processing algorithm (Art, Gnanadesikan and Kettenring 1982). *K*-means clustering (MacQueen 1967) is used to dissect these data in the range from two to ten clusters. Most of the 50 runs are similar, and the cluster analysis is simply splitting off a few outliers as the number of clusters increases. Specifically, for the 50 runs with the ZI retailer as moderator, 45 of the 50 runs are the same. For the runs with the naive retailer, 49 of the 50 runs are the same. And for the sophisticated retailer, while this shows less homogeneity than the other retailers, 42 of the 50 runs are the same. We report only the majority case for each retailer, discarding the other runs as atypical. Other outcomes we present later include the total category volumes and brand profits associated with these solutions, as well as an analysis of promotions offered to retailers and their choice of whose promotions to place in their stores. Finally, we analyzed the historical actions of brands and retailers. We did this by clustering shelf prices using the stylized actions shown in Table 3 as cluster seeds which, combined with the promotional

analysis we discussed earlier, provides an equivalent picture of history. For reasons of sample size, especially for promotions, here we combine the four brands, rather than take them brand by brand.

## VI. RESULTS

### *A. Actions of the brand-agents*

Table 5 shows the stylized actions of the brand-agents, and their historical equivalents. As in our previous work the brand-agents promote more than do real brand managers. There can be many reasons for this result that are yet to be incorporated in our model. Brand managers may face constraints which we do not model (budgets, practicalities, coordination across product categories, et cetera), or may have concerns for the longer-term effects of over-promoting their brands. We build the consumer response model from a 78-week database and thus cannot incorporate any longer-term trends in price sensitivity. But the striking feature of Table 5 is that when a sophisticated retailer moderates promotions, the outcome is much closer to the historical pattern, indeed the closest we have come in our program of research. Thus the form of retailer has a major role in controlling the promotional behavior of the brands.

[Table 5 about here]

The differences between the ZI retailer and the naive retailer are less striking but still obvious. In particular, the *Maxwell House*, *CFON* and *Hills Brothers* brand-agents change their behavior between the ZI and naive retailer cases. *Maxwell House* and *Hills* decrease their use of the Low2 action, whereas *CFON* decreases its use of the HIGH2 action. With the results for the sophisticated retailer case, this shows that our brand agents are optimized to the particular environment they face. In all three cases the frequency of actions differs between brands. This shows that the agents are also reflecting the asymmetries of their market position and cost structures.

### *B. Category volumes and brand profits*

Table 6 displays the total category volumes and brand profits associated with the three cases. Again, the obvious difference is between the results for the sophisticated retailer and the ZI and naive retailers. Indeed, the results for the latter two do not differ much in magnitude. Taking total category volume first and running some simple *t*-tests (with the two-tailed probability levels corrected for multiple tests) reveals

the following. There is no significant difference between the total category volumes achieved by the ZI retailer versus the naive retailer, but both are significantly *greater* than that achieved by the sophisticated retailer ( $p < 0.002$ ). *Maxwell House* and total (four-brand) profits show a similar pattern. Again there is no significant difference between the ZI retailer and the naive retailer, but both profit levels are significantly *greater* than for the sophisticated retailer ( $p < 0.002$ ). The results for the other brands are more complex statistically. *CFON* has a U-shaped pattern, where profits are greatest for the naive retailer and significantly different from the ZI and sophisticated case ( $p < 0.002$ ). However, the latter are not statistically different from each other. For *Hills Brothers*, all cases are significantly different from one another ( $p < 0.002$ ), with the sophisticated case having the lowest value. However, by magnitude the ZI and naive retailer values are similar. Perhaps most interestingly of all, the market leader, *Folgers*, shows no significant differences between the three cases. Its profits are always the greatest and seemingly not affected by different retail moderators, even though total category volume and total (four-brand) profits clearly fall for the sophisticated case.

### *C. Retailer choice and promotional competition*

The final set of results concerns the choices of the retailers. Which brands offered promotions at the same time (which we term “clashes”) and whom does the retailer prefer? Table 7 shows the simulation results. We omit the ZI case since there clashes are resolved by coin tossing, with no attention paid to brand or promotion. In the Table we only show two-way clashes, which are the main instance. Three- and four-way clashes are less common and are discussed below. For the naïve-retailer case there are 30 two-way clashes and 18 three-way clashes, but no four-way clashes. In the two-way clashes it is notable that the leader, *Folgers*, does not win most of its promotional contests, even against the weakest of the four brands, *Hills Brothers*. Similarly for three-way clashes, the best *Folgers* does is to match *Hills Brothers* in clashes between these two brands and *Maxwell House*. But of course the naive retailer is blind to brand. The other notable feature of the naive retailer clashes is that nearly every week of the 50 weeks involves a clash. Only twice did a brand offer a promotion without competition.

For the sophisticated retailer case, there are 16 two-way clashes, only 1 three-way clash and again no four-way clashes. And there are 23 occasions in which only a single brand offered a promotion, with no competition from other brands. This is a striking result in that not only does the sophisticated moderation lessen promotional competition, but the brand-agents have also evolved to avoid clashes. Where two-way clashes do occur, *Folgers* does much better than in the naive case because the sophisticated retailer recognizes its asymmetric power. Nonetheless, *Folgers* does not always win the competition: sometimes the other brand offers a more compelling promotion. Indeed, even the weakest brand, *Hills Brothers*, wins some competitions, as can be seen in Table 7. The reasons for these differences between the naive and sophisticated cases are as follows.

Each brand has three actions (LOW1, LOW2, and Low1 in Table 3) recognized as promotions. But for any brand, the three are not equally attractive to the sophisticated retailer in terms of category volume because of brand asymmetries. The rankings we use to resolve clashes vary with the retailer model. For the naive retailer, the descending ranking is simply: LOW1, LOW2, and Low1, since the naive retailer is brand-blind. Any clashes where two brands propose the same promotion are resolved randomly. The sophisticated retailer, however, has a more nuanced descending ranking: *Folgers*-Low1, *Maxwell House*-LOW1, *Maxwell House*-LOW2, *CFON*-LOW1, *CFON*-LOW2, *Folgers*-LOW1, *Maxwell House*-Low1, *CFON*-Low1, *Folgers*-LOW2, and then any promotion by *Hills Brothers*, from Table 4. That is, although the proposed actions are identical across brands, the sophisticated retailer recognizes that the brand asymmetries mean an identical action will result in different outcomes across brands, and ranks the proposals accordingly. Thus this retailer can always choose between proposals. The numbers in Table 7 are the realized frequencies of clashes (and outcomes) which occurred in the simulations.

## VII. DISCUSSION AND CONCLUSIONS

### A. On results and hypotheses

Our results have presented us with four surprises. First, category volume and brand profits are lower with the sophisticated retailer than with the naive retailer model. Second, except, that is, for *Folgers*, the dominant brand, whose profits are preserved in an absolute sense and rise with sophistication if viewed as

share-of-total profits (Table 6). Third, despite this (or perhaps because of it), sophisticated moderation results in brand behavior which is the most realistic of the three models of moderation (Table 5). Finally, with sophisticated moderation, when the retailer looks at which brand is proposing a promotion, as well as the promotional market actions (Table 4), the brand-agents learn (in an evolutionary sense) to avoid promotional clashes.

Should we conclude that it is difficult to justify the costs (whatever they are) of sophisticated moderation by improved category volume? It is not true that naive moderation strictly dominates sophisticated moderation, since *Folger's* preserves its profits with sophistication, and *Folger's* is the dominant brand. *Folger's* is well placed to compensate the retailer for any decrease in category volume that results from sophisticated moderation. Whether any such side-payments occurred is moot, and it is also possible the spill-over effects of promoting the dominant brand on sales in other product categories compensate the sophisticated retailer. Few consumers shop for just one item. And of course, this moderation reduces the profits of *Folger's* competitors, which has advantages for retaining dominance.

That our brand-agents' behavior is most realistic with sophisticated moderation suggests that real-world moderation is likely more sophisticated (that is, brand-aware) than less. That our brand-agents clash less often with sophisticated moderation than with the other cases shows that the evolutionary process is learning to avoid the risk of losing. Losing means having to price at an uncompetitive shelf price.

We conclude that H1 receives no support: ZI moderation equals naive moderation. Further, although naive moderation provides higher category volume for the retailer than does sophisticated moderation, the behavior of the brand-agents in the simulations suggests that the dominant brand induces brand-awareness in the retailer's decision-making. H2 is thus not supported in the sense of category volume, although future research may reveal that side-payments or spin-offs make it more profitable for the retailer overall.

*B. On the power of evolutionary computing to explore complex economic phenomena*

Evolutionary algorithms were first developed in an engineering tradition of design and optimization. This is true of Holland (1975), in the monograph introducing biological models (the GA) of adaptation in psychology, economics, control theory, artificial intelligence, sampling, and in computer sciences.

“Basically, adaptive processes are optimization processes” (p.v). This emphasis on the GA to seek better solutions in a complex search space has continued as social scientists apply evolutionary algorithms to their problems. And often the underlying theory assumes that agents are optimizing. That is also true of our agents, who try to maximize their profits.

But we are not seeking to optimize marketing strategy for brands or supermarkets: we are not engineers. Instead, we are using the tools of evolutionary computation to explore the real-world behaviors of brand managers and retailers that moderate competition among brands. To do this, our model of their behavior makes the reasonable assumption that these real-world managers are not behaving randomly, but are purposely making more or less informed decisions. That is, we are using evolutionary algorithms to explore real-world phenomena, not seeking to improve on reality.

Of course, to improve on managerial practices it is first necessary to understand them, so we might see exploring these phenomena as a precondition to suggesting improvements, but that is not our goal here. As we stated in the beginning, we lack full understanding of the interactions among oligopolists, and introducing a moderator of their rivalrous dance — the retailer — makes such understanding more difficult to achieve. We hope we have advanced that understanding a little by showing that the way in which this retailer moderates is critical to the dance that emerges

### *C. Next steps*

Earlier we mentioned influences on the brands and retailers that we have not tried to model here. An obvious extension of this work would be to develop a similar model for a different supermarket chain, and to allow competition between the two chains for patronage. We could assume independent proposals for promotion from the same brand across the two retailer chains, or, more interestingly, we could link a brand’s proposals in one chain to its proposals in the other chain. The chains too would be asymmetric in

their costs and their market appeal. When we use evolutionary models to improve understanding of real phenomena, we must model the asymmetries we see in the real world, and not assume homogeneity.

The other obvious extension from our findings is to understand the impact of promotions in one category on other categories in the supermarket. That is, to model the total impact of a promotion on the retailer's profits. This however, either needs more extensive datasets than are currently available or developing a simpler, perhaps stylized, model of other categories from more limited data. Perhaps it is possible to derive a simple model of total impact without knowing the specifics of other categories?

Finally, and much more ambitiously, we can replace both the econometric models of both retailers and consumers with bit-string agents. For retailers, this would be a straightforward extension of the brand-agents, since retailers could also be assumed to be profit-maximizing. For consumers, however, this is a more difficult problem because it is harder to define a fitness score around their consumption satisfaction or utility.

#### REFERENCES

- Art, D., Gnanadesikan, R. and Kettenring, R. (1982), "Data-based Metrics for Cluster Analysis," *Utilitas Mathematica*, 21A, 75-99.
- Axelrod, R. (1987), "The evolution of strategies in the iterated Prisoner's Dilemma," in: L. Davis, ed., *Genetic Algorithms and Simulated Annealing* (Pittman, London) 32-41.
- Carpenter, G., Cooper, L., Hanssens, D. and D. Midgley (1988), "Asymmetric market share models," *Marketing Science*, 7(4), Fall, 393-412.
- Cliff, D. (2001) "Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets," Hewlett-Packard Technical Report HPL-2001-99.
- Cliff, D., and J. Bruten (1997), Zero is Not Enough: On The Lower Limit of Agent Intelligence For Continuous Double Auction Markets, HPL-97-141, November.
- Cooper, L. and M. Nakanishi (1993). *Market share analysis*. Boston, MA: Kluwer.

- Das, R., J.E. Hanson, J.O. Kephart and G. Tesauro (2001), "Agent-human interactions in the continuous double auction," in: B. Nebel, ed., *Proceedings of the 17th International Joint Conferences on Artificial Intelligence (IJCAI)*, Seattle (Morgan Kaufmann, San Francisco) 1169–1187.
- De Jong, K.A. (1975), An analysis of the behavior of a class of genetic adaptive systems. PhD dissertation.
- Duffy, J. (2006) Agent-Based Models and Human Subject Experiments, in the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by L. Tesfatsion and K. L. Judd, in the series *Handbooks in Economics*, (general editors K. Arrow and M. D. Intriligator), (Amsterdam: Elsevier Science), pp. 949-1011
- Fudenberg, D. and Levine, D.K. (1998) *The Theory of Learning in Games*, MIT Press.
- Gathercole, C. and P. Ross (1997) Small Populations over Many Generations can beat Large Populations over Few Generations in Genetic Programming, *Genetic Programming 1997: Proceedings of the Second Annual Conference*, Stanford, Morgan Kaufmann Publishers Inc. San Francisco, pp.111-118.
- Gode, D. and S. Sunder (1993), "Allocation efficiency of markets with Zero Intelligence traders: market as a partial substitute for individual rationality," *Journal of Political Economy* 101: 119-137.
- Grefenstette, J. (1986), Optimization of control parameters for genetic algorithms, *IEEE Transactions on Systems, Man and Cybernetics*, 16 (1), 122 – 128, Jan-Feb .
- Holland, J.H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor: University of Michigan Press.
- MacQueen, J.B. (1967), "Some Methods for Classification and Analysis Of Multivariate Observations," *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1, 281-297.
- Marks, R.E. (1992) Breeding optimal strategies: optimal behavior for oligopolists, *Journal of Evolutionary Economics*, 2, 17-38.

Marks R.E., Midgley D.F., Cooper L.G. (2006), Co-evolving better strategies in oligopolistic price wars, in *Handbook of Research on Nature-Inspired Computing for Economy and Management*, edited by Jean-Philippe Rennard, (Hershey, PA: Idea Group Inc.), pp. 806-821.

Midgley, D.F., Marks, R.E., and Cooper, L.G. (1997) Breeding competitive strategies, *Management Science*, 43(3): 257-275, March.

Mitchell, M., (1996), *An Introduction to Genetic Algorithms*, MIT Press.

Schaffer, J. D., Richard A. Caruana, Larry J. Eshelman, Rajarshi Das (1989) A study of control parameters affecting online performance of genetic algorithms for function optimization, *Proceedings of the Third International Conference on Genetic algorithms*, Morgan Kaufmann Publishers Inc. San Francisco, pp. 51-60.

Schraudolph, N.N. & J.J. Grefenstette, *A User's Guide to GAUCSD 1.4*, Technical Report CS92-249, CSE Department, UC San Diego, 1992 , at <http://www.agsm.edu.au/~bobm/teaching/SimSS/GAUCSD.pdf>

Van Valen L. (1973): "A New Evolutionary Law", *Evolutionary Theory* 1, 1-30.

Vriend, N. (2000), "An illustration of the essential difference between individual and social learning and its consequences for computational analyses," *Journal of Economic Dynamics and Control* 24: 1–19.

**TABLE 1: OLS regression coefficients for the Naive Retailer**

<b>Variable</b>	<b>Coefficient</b>	<b>t-value</b>	<b>p value</b>
Intercept	1491.3	17.3	<0.001
Discount from reference shelf price	1293.1	3.1	<0.01
Proportion of stores with a major promotion	905.6	3.0	<0.01
Proportion of stores with a minor promotion	948.6	2.0	<0.05

**TABLE 2: Regression coefficients for the sophisticated retailer**

Variable	Estimate	t-value	p value
Intercept	6.43	7.9	<0.001
Reference price (average price of non-promoted brands)	-1.16	-4.7	<0.001
Lag 1 parameter	0.35	3.0	<0.01
<b>Major Promotions</b>			
Folgers	0.94	8.2	<0.001
Maxwell House	1.33	15.9	<0.001
CFON	1.24	11.2	<0.001
Hills Brothers	0.88	6.3	<0.001
<b>Minor Promotions</b>			
Folgers	2.04	8.1	<0.001
Maxwell House	1.43	3.1	<0.01
CFON does not run minor promotions in this chain			
Hills Bros	0.00		NS
<b>Impact of other supermarket chain</b>			
Log (category volume in other chain)	0.38	5.4	<0.001
Reference price in other chain	0.52	5.3	<0.001
Major promotion in other chain	-0.22	-1.8	< 0.04 One tailed
Minor promotion in other chain			NS

Shaded coefficients are those relevant to a choice between competing promotions

**TABLE 3: Stylized actions**

<b>Action label</b>	<b>Price (\$)</b>	<b>Feature (% of stores)</b>	<b>Display (% of stores)</b>
LOW1	1.89	85	85
LOW2	1.99	95	95
Low1	2.09	0	80
Low2	2.19	0	0
High1	2.39	0	0
High2	2.49	0	0
HIGH1	2.55	0	0
HIGH2	2.69	0	0

Shaded promotions subject to store policy of only one allowed per week

**TABLE 4: Promotional category volumes for naive and sophisticated retailers****(A) Naive Retailer**

Type of Promotion	Category Volume
LOW1	3115
LOW2	3076
low1	2845
Reference volume without promotion	1491

**(B) Sophisticated Retailer**

Brand	Type of Promotion	Category Volume
<i>Folgers</i>	low1	3265
<i>Maxwell House</i>	LOW1	3250
<i>Maxwell House</i>	LOW2	3127
<i>CFON</i>	LOW1	3090
<i>CFON</i>	LOW2	2981
<i>Folgers</i>	LOW1	2612
<i>Maxwell House</i>	low1	2609
<i>CFON</i>	low1*	2609
<i>Folgers</i>	LOW2	2541
<i>Hills Brothers</i>	LOW1	2525
<i>Hills Brothers</i>	LOW2	2462
<i>Hills Brothers</i>	low1	1587
Reference volume without promotion		1541

\* CFON does not run display-only promotions in the period we observe. To estimate what the impact of such a promotion might be we use the coefficient from an equivalent brand, *Maxwell House*.

**TABLE 5: Stylized actions from Monte Carlo simulations (Percentage of actions)**

**Stylized Actions**

	<b>LOW1</b>	<b>LOW2</b>	<b>Low1</b>	<b>Low2</b>	<b>High1</b>	<b>High2</b>	<b>HIGH1</b>	<b>HIGH2</b>
<b>Zero-Intelligence Retailer</b>								
<i>Folgers</i>	23	0	0	22	1	54	0	1
<i>Maxwell</i>	24	0	0	20	2	54	1	0
<i>CFON</i>	20	0	0	0	0	56	2	22
<i>Hills</i>	33	0	1	16	0	49	1	0
<b>Naive Retailer</b>								
<i>Folgers</i>	21	0	1	20	1	54	0	1
<i>Maxwell</i>	25	1	2	4	2	62	3	2
<i>CFON</i>	22	1	1	2	2	59	5	8
<i>Hills</i>	25	1	1	4	0	67	1	2
<b>Sophisticated Retailer</b>								
<i>Folgers</i>	8	7	11	15	8	32	9	11
<i>Maxwell</i>	5	4	7	15	13	32	12	12
<i>CFON</i>	6	6	6	10	12	30	14	15
<i>Hills</i>	5	6	8	12	9	32	12	16
<b>Historical</b>								
4 brand data	2	7	2	13	24	12	33	7

**TABLE 6: Total category volumes and brand profits**

	<b>Zero-intelligence retailer</b>		<b>Naive retailer</b>		<b>Sophisticated retailer</b>	
	<b>Mean</b>	<b>Std.dev</b>	<b>Mean</b>	<b>Std.dev</b>	<b>Mean</b>	<b>Std.dev</b>
Total Category Volume	1756	17	1763	40	1617	63
<b>Profits</b>						
Folgers	331	11	324	17	330	26
Maxwell House	244	12	241	13	203	14
CFON	66	3	76	10	63	12
Hills Brothers	45	1	42	3	29	3
Total (Four Brand)	686	8	682	17	625	28

**TABLE 7: Promotional Clashes**

**Two-way clashes for the naive retailer case (30)**

	Folgers	Maxwell	CFON	Hills
Folgers		2 to 1 Maxwell		5 to 4 Hills
Maxwell			8 to 3 Maxwell	
CFON				4 to 3 CFON
Hills				

**Two-way clashes for the sophisticated retailer case (16)**

	Folgers	Maxwell	CFON	Hills
Folgers		2 to 1 Folgers	1 to 1	2 to 1 Folgers
Maxwell			1 to 1	1 to 1
CFON				2 to 2
Hills				