Refining the breeding of hybrid strategies

By

Robert E. Marks

David F. Midgley

Australian Graduate School of Management

University of New South Wales

Sydney NSW 2052 AUSTRALIA

Phone +61-2-9931-9200

Fax +61-2-9313-7279

bobm@agsm.unsw.edu.au

davidm@agsm.unsw.edu.au

and

Lee G. Cooper

Anderson Graduate School of Management

University of California Los Angeles

Los Angeles CA 90095 USA

lcooper@agsm.ucla.edu

Current draft: October 5, 1998

ABSTRACT

We extend our earlier work using artificial agents to model multi-period games between competing brands in an oligopoly. We do so by developing a multiple-population genetic algorithm in order to allow customized agents for each brand. We also consider more competitors and more possible pricing actions per competitor than before, and we evaluate the robustness of our results by Monte Carlo methods. All these developments have been facilitated by writing better code and by increases in computing power since our original work. We find that:

- Customized agents outperform traditional genetic algorithm approaches;
- The addition of a niche player changes the behavior of the major brands; and
- An increase from 4 to 8 actions per agent results in more capable agents.

In addition, we report on two surprising effects. The *Holyfield-Tyson effect:* whereby sophisticated agents do not perform that well against primitive agents; and the *Frankenstein effect:* whereby agents developed in competition with other agents exhibit different behaviors when competing with the historical actions of managers.

Overall, we believe the strength of our approach results from the use of an empirically grounded fitness function with which to test our assumptions and approaches.

Introduction

We are interested 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 situation faced by managers for many product categories sold in supermarkets. Through a substantial literature built on the analysis of scanner data, much is known about consumer response in these situations, and methods for modeling this response are well-established (Cooper and Nakanishi 1988). Less is known about the competitive actions of managers and there have been few attempts to model the repeated game evident in these product categories.

In earlier work (Midgley, Marks and Cooper 1997), we showed how the actions of each brand manager could be modeled as an outcome of finite automata playing a repeated game. Using the Axelrod and Forrest representation (Axelrod 1987) of these artificially adaptive agents as bit strings, we employed a genetic algorithm and a market-response model to co-evolve artificial agents for brands of coffee in a regional U.S. market.

These agents were specified as partitioning the previous actions of competitors into a small number of states and selecting an action that would be profitable for them in the next period of the game from a similarly small number of available actions. This process can be thought of as defining perceptions of possible states of the market and developing mappings from these perceptions to an action for the next period.

In our earlier work, agents were limited to a set of four actions and their perceptions restricted to four equivalent states for a single previous period of the game. This resulted in agents whose mapping from perceptions to actions could be represented by 134 bits. The genetic algorithm (GA) was then used to evolve mappings that maximized profits over a number of multi-period games in which various agents were pitted against their competitors. The GA achieved this by selecting and *hybridizing* the better performing strategies from each generation of these games to create more profitable strategies for the next generation, in an iterative cycle of selection, mating and improvement. Profits were computed from a market response model that estimates brand sales, given the actions of the competing brands.

These hybridized agents performed well both in these games and when a single agent was pitted against the historical actions of human brand managers. The latter test is “unfair”— as there is no opportunity for the human managers to respond to the agent— but this was a surprising result given the simple agents we used. Moreover, in developing these agents, we learned that store policies and demand saturation were also important to the realism of our results. Store policies are important because these can act to constrain the frequency of temporary price reductions and promotional displays, and demand saturation is important because there is a limit to the amount of any product that can be consumed or stored in a given period. Both these constraints set the environment within which managers and agents compete.

While these results are encouraging, we recognize that the agents and procedures we use are capable of further development. There are two important limitations to our earlier work, the first methodological and the second concerning the sophistication of our agents.

First, the GA used a single population of 25 artificial agents, scoring the profitability of each string differently, depending on which of the three brands the agent was designated as in a particular simulation game. This is analogous to training all brand managers in the same business school. It was done because the GA software available at the time only addressed the single-population case. But in our situation—where consumer response and firm costs differ by brand—it would be desirable to have a multi-population GA. There are also issues around the robustness of our results. In our previous work we only ran one simulation per experiment whereas Monte Carlo simulations from different starting conditions would be a better methodology. Advances in hardware now make such multiple simulations possible.

Second, we modeled only the three main players in a market that has nine brands and we only allowed our agents four possible actions, while human managers used a far greater number. These choices were made because the high computing demands of GA applications made it difficult to complete more complex simulations. In the intervening period, hardware has improved and we have learned to make GA software more efficient. This permits more realistic representations of brand management to be modeled.

This paper reports the consequences of relaxing these restrictions. In particular, how the performance of the agents improves as we allow brand-specific responses with separate GA populations, as we co-evolve four artificial agents instead of three, and as we use a set of eight possible actions instead of four. The paper also begins the task of modeling the learning inherent in the actions of human managers. While these managers use more actions than do our agents, the actions they use are only a subset of all possible actions. Just how they came to select this subset, and how optimal it is compared with other possibilities, are important research questions. By examining agents at different stages of their 'evolution' and by using different sets of actions, we can begin to address these issues.

The structure of the paper is as follows. We first describe necessary improvements to our modeling procedures that speed up optimization, assess the robustness of our results, and allow multiple populations. We then present the results from a series of experiments: Experiment 1 where we examine the impact of multiple populations; Experiment 2 where we increase the number of players from 3 to 4; and Experiment 3 where we increase the number of possible actions from 4 to 8. Finally, we examine co-evolution and genetic drift in Experiment 4, and the nature of managerial learning in Experiment 5. We conclude by identifying issues of concern and areas for future research.

Methodology

We have made a number of improvements to the basic methodology described in Midgley, Marks and Cooper (1997). These improvements relate to Monte Carlo methods and multiple-population genetic algorithms.

Monte Carlo simulations

In order to assess the robustness of our results, we now perform Monte Carlo simulations for each of our experiments rather than the single simulations of our previous paper. These multiple runs are easily achieved, as the GA needs a random number seed to generate its initial population of strings. We simply start each Monte Carlo simulation with a different seed. It should be noted, however, that Monte Carlo methods have been greatly facilitated by the steps we have taken to speed up simulations—together with increases in computing power. Previously, running repeated simulations would have taken too long. In particular, we now prescreen the strings in the initial population to cull illegal genotypes—that is, strings that would violate of store policy. This "filtering" greatly accelerates agent learning. We have observed convergence occurring from the 20th generation with filtering compared with the 70th generation previously—a dramatic improvement in performance.

Identifying patterns of competitive interaction. The Monte Carlo simulations give us confidence in our results but require us to summarize large amounts of data on the behavior of the agents. We have developed a methodology for this.

As the bit strings evolve over the generations, they typically begin to favor some actions at the expense of others. In other words, they learn which actions maximize profits, given the constraints and the behavior of their competitors. Over Monte Carlo simulations, different patterns of actions can potentially emerge, as the populations of agents evolve from random starting seeds. These different patterns manifest themselves in the relative frequency with which the agents use their assigned actions as they compete. To obtain frequencies based on fully optimized agents, we use the final generation of the simulation. We then cluster analyze these frequencies to identify patterns. This cluster analysis is done jointly for all the competing brands, rather than looking at each brand in isolation, because the actions of a brand in any one Monte Carlo run are dependent on the actions of its competitors in that same run. If we find similarities in these patterns of interaction across simulations, then we can be confident that our results are robust to random changes in the starting seeds.

Thus cluster analysis is done to summarize the data rather than to identify a 'true' number of clusters—an approach that Everitt (1980) calls "dissection". We used a k-means algorithm in a standardized procedure to generate ten patterns of competition from the 50 Monte Carlo simulations run for each experiment. We chose the number ten to clearly identify differences, should they exist. We visually inspect the patterns to see how different they are from each other, and we label them to describe the strategy represented by each (e.g. 'Every Day Low Price'). In this manner we can simplify a large volume of results.

Multiple-population simulations

Our earlier work—in common with many published uses of GAs—relied on a single population of strings (agents). We handled the different brands by using different payoff matrices for each (computed from their distinct market response and costs) and by applying each string from the common population to each payoff matrix in turn. Thus in a three-brand simulation, each string would generate three different profit outcomes. This meant that the profitability of a string varied according to situation, and hence the GA itself was subjected to greater variance in its search for higher-performing strings. So long as the profit surfaces of the three strategic brands had a similar topology, then this variance should not have created much difficulty for the GA. The single population would behave as though the payoffs were noisy, but not pathological, in the way that opposite slopes of profit functions might induce.

Nonetheless, we determined to develop a simulation with multiple populations. This would allow us to reflect the differing responses brands invoke from consumers, and the differing costs they face, more accurately in our simulations.1 Another compelling reason for developing multiple populations relates to the sharing of genetic information. With a common population, all brands share the same genetic information and develop similar strategies. While it is true that managers may change companies within an industry, and thus share strategies between firms, we doubt that this occurs to the same extent as with our single-population agents. Distinct populations allow the agents to develop differentiated strategies, and we believe that this is closer to the competitive realities of these markets.

Developing a multiple-population GA was not a trivial exercise, since we have three or four different players competing at the same time in each period, and necessitated getting under the hood of our GA engine, the UCSD version of GENESIS. The net effect of G. M. Shiraz's reprogramming is that the simulations are much faster: by using all the information generated in each interaction of players. Indeed, with three populations of players, the new code is almost as fast per trial as the old code was with a single population.

Because of the stochastic nature of the simulations, we have performed some Monte Carlo simulations to compare the convergence and outcomes of moving from a single population with three brands to three distinct populations, one per brand. These results are described under Experiment 1 below. Having made these methodological improvements we now turn to developing more sophisticated and realistic simulations.

Four strategic players

A natural extension of the earlier work has been to increase the number of strategic players. With the new multi-population code, it has been relatively easy to extend the simulations to a fourth player, at some cost in the complexity of the bit-string chromosomes. These strings grow in length from the 134 bits in our previous paper to 12,312 bits for the most complex experiment reported in this paper (one-week memory and eight possible actions per player, including bits for the phantom memory of the first week's play2).

The choice of the fourth player is not obvious. The fourth-ranked brand by market share—*Master Blend*—belongs to the same owner as one of the three main brands and was relatively inactive in terms of price promotions for the period we have data. Instead we chose a smaller player—*Hills Bros*.—that is known to be a strong competitor for some of the major brands. This choice seemed to us to add a more interesting dimension to the competitive game. These results are reported under Experiment 2 below.

Eight possible actions per player

In our earlier work we chose four as the number of possible actions per player and we chose the values for these actions from those commonly used by managers. These values differ somewhat by brand—a natural consequence of the differentiating strategies used by managers. They are also a combination of price, advertised feature and store display. "Featuring" involves the store promoting the brand in local newspapers—for which they charge the manufacturer and which normally invokes a stronger response from consumers than does a price cut alone or a store display. A featured low price is thus the type of promotion that is controlled by store policy and incorporated into our constraints. 3

Using four actions meant that the artificial agents were more constrained than their historical counterparts had been, and thus we denied them the opportunity to learn what the brand managers must have learnt through experience and corporate memory: the boundaries of extreme behavior. By increasing the number of possible actions to eight, we hoped to give our agents the opportunity to demonstrate that the four actions used earlier were robust, and that our assumption of a mature oligopoly was therefore correct.

Eight possible actions meant an increase in the complexity of the simulations because of the longer bit strings required. As before, we identified the appropriate actions for each brand by a process of cluster analysis and visual inspection of historical data. We also used historical data from three chains rather than simply from the one chain that is the focus for our simulations. This provided us with a broad sample of possible actions. Six actions for each brand were chosen by this process—to which we added the highest price and the deepest price promotion observed for the brand in the period. The final eight actions allow the agent to have a wider-ranging and more varied set of actions than in our earlier work. The results of simulations based on eight actions are discussed under Experiment 3 below.

Co-evolution and genetic drift

The artificial agents learn through application of the recombinant evolutionary techniques of the GA. This is clear when the agents are solutions to a static problem, as has been the most usual application of GA techniques in say, engineering. But Marks (1992) and others following have bred agents against each other, a process that biologists term “co-evolution”.4 Against a static environment, improving fitness scores readily reveals the progress of the artificial agents, but against a dynamic environment comprised of like artificial agents, scores may not rise from generation to generation.

In our earlier work we attempted to show the competence of our artificial agents by pitting them against the historical actions of managers, but some criticism has been made that this overstates the skills of the artificial agents and understates the skills of the managers. That is, in this open-loop setting, the managers have no opportunity to respond to the actions of artificial agents, as their plays are given and unchanging. Here we attempt to show how the agents have learnt by taking those evolved from 100 generations and playing them not against the frozen patterns of their historical opponents, but rather against agents evolved after only eight generations. We term this process pitting “sophisticated” agents against “primitive” agents. The results of these competitions are given under Experiment 4 below.

Managerial learning

In our final experiment—Experiment 5—we investigate an analogous learning process, that of the brand managers. We did this by contrasting the profitability and behavior of agents using the eight actions derived from historical data with agents using eight randomly determined actions. The actions we see in the historical data are the end result of decades of competition and managerial learning. By contrasting agents using these actions with those using random actions, we attempt an assessment of the consequences of this learning.

Results

Experiment 1 — multi-population simulations

We did 50 Monte Carlo runs under a common-population scenario, and 50 runs under a distinct—brand specific—population scenario. We used the three-brand, four-action model from our previous paper but with the methodological improvements noted above. Figure 1 shows how the 'industry profits' evolve during the genetic optimization.

Figure 1 about here

The industry profits shown in Figure 1 are the averages observed across the 50 runs under the common- and distinct-population scenarios. For the common-population scenario these are computed from 50 simulations, each with a common population of 25 agents; for the distinct-population scenario these are computed from 50 simulations, each with three distinct populations of 25 agents. The distinct-population scenario generates higher profits and converges more rapidly than does the common-population scenario. This is because the latter is a noisier environment for optimization and because there are benefits to developing a distinct solution for each brand. While in aggregate these benefits are not particularly large—increasing industry profits by only 4%—if we examine brand profits, we see a significant reallocation. When given their own populations of agents, *Folgers* increases its profits by 3%, and *Maxwell House* increases its profits by 24%, but *Chock Full O'Nuts* loses some 16% of its profits. Distinct populations allow the agents for *Maxwell House* to better capitalize on that brand's strengths, to produce more effective competitive behavior.

How then does the behavior of agents evolved under the distinct-population scenario compare with those evolved under our earlier common-population scenario? We applied our

Figure 1 Evolution of Industry Profits with Common and Distinct Populations

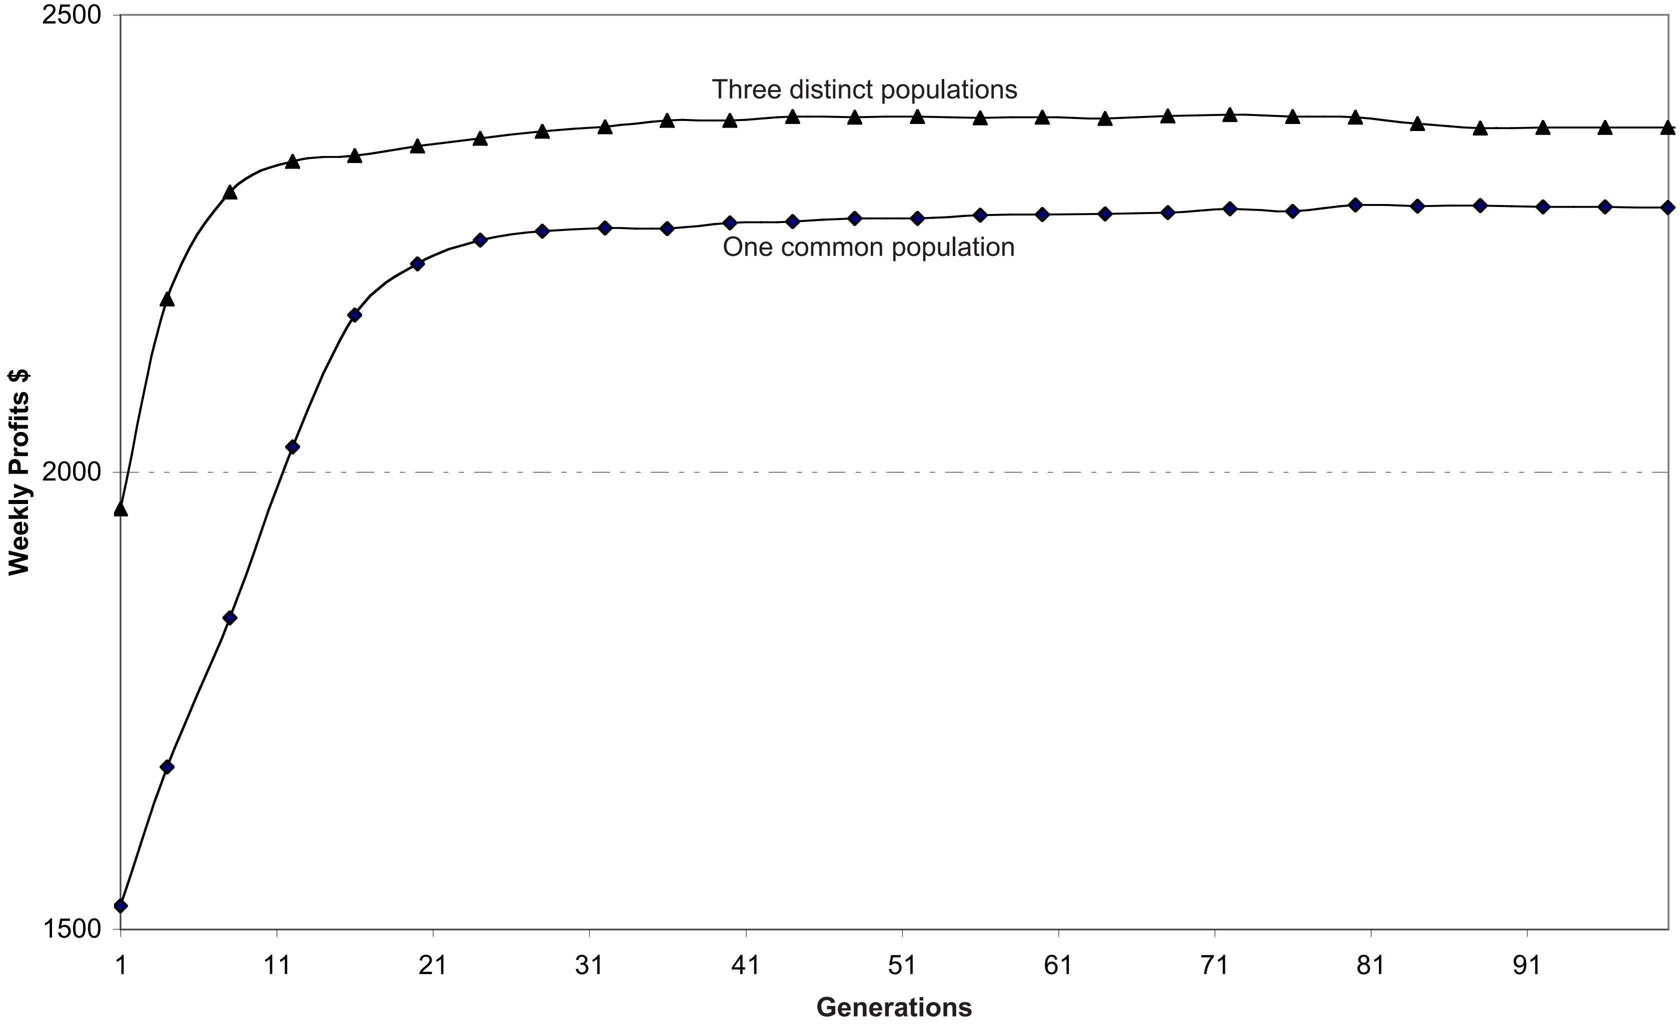

cluster dissection methodology to the action frequencies from the Monte Carlo simulations. Table 1 shows the results for the one-population, four-action scenario and Table 2 the equivalent results for three populations.

Tables 1 and 2 about here

In Table 1 we find that the 50 simulations are very homogenous. Patterns 1 and 2 are almost identical and represent 32 of the 50 simulations. In fact, detailed inspection of the remaining patterns suggests that 48 simulations are essentially the same, one simulation produces low performance and that only Pattern 3 represents a high-performing alternative pattern of competition. The common pattern of interaction is thus EDLP for *Folgers* and *Chock Full O' Nuts* and Wide Pulsing for *Maxwell House*. Note that these represent a linked set of strategies—they co-evolved over the 100 generations. This discussion also demonstrates how we dissect the data from the Monte Carlo runs and assign labels to the strategies.

Table 2 also shows homogeneous patterns of competition. Patterns 1 and 2 are almost identical and account for 41 of 50 simulations. Inspection of the remaining patterns reveals that all except Pattern 3 have worse performance than those tabled. And Pattern 3 only differs from 1 and 2 in the behavior of *Chock Full O'Nuts*.

If we contrast the one-population, four-action experiment with the three-population, four-action experiment, the basic strategies remain the same. This result is not unsurprising, as the actions available to the agents remain the same. But, while the basic strategies remain the same, the detailed profiles of each change somewhat, and this allows the distinct-population agents to differentiate themselves more. For example, the Wide Pulsing strategy of *Maxwell House* involves 47% of its actions being at the lowest price in the distinct-population experiment, versus 32% in the common-population experiment. Indeed, the main beneficiary of allowing these agents to respond more directly to the individual characteristics of the brands is the market leader, *Maxwell House*.

Overall we conclude that moving to distinct populations has resulted in higher-performing strings. Distinct populations also result in greater heterogeneity in the performance of the agents for each brand. For example, the coefficients of variation of brand performance for the common-population case are 3%, 5% and 2% respectively, and for the distinct-

Table 1: Patterns of competition among evolved agents—common population and four actions

	Actions Low price			High price		
Pattern 1, 21 runs[a]	**1**	**2**	**3**	**4**	**Average Profit**	**Strategy**
Folgers	1[b,c]	**98**	0	1	$1,022	EDLP
Maxwell House	**32**	7	14	**47**	$631	Wide Pulsing
Chock Full O' Nuts	0	**100**	0	0	$633	EDLP
Pattern 2, 11 runs	**1**	**2**	**3**	**4**	**Average Profit**	**Strategy**
Folgers	0	**97**	2	1	$1,011	EDLP
Maxwell House	**33**	4	10	**53**	$625	Wide Pulsing
Chock Full O' Nuts	0	**98**	0	2	$630	EDLP
Pattern 3, 1 run[d]	**1**	**2**	**3**	**4**	**Average Profit**	**Strategy**
Folgers	**46**	**52**	0	2	$1,082	Promote to the Max
Maxwell House	**30**	0	**34**	**36**	$623	Wide Pulsing
Chock Full O' Nuts	0	**50**	0	**50**	$707	Pulsing on Shelf Price

a: patterns of competition are computed during the 100th generation from all combinations of 25 agents playing 52-week games and from 50 Monte Carlo simulations.

b: row percentages, key numbers in bold.

c: shaded areas identify the actions constrained by store policy.

d: best performing of remaining patterns.

Table 2: Patterns of competition among evolved agents—three distinct populations and four actions

	Actions Low price			High price		
Pattern 1, 25 runs[a]	**1**	**2**	**3**	**4**	**Average Profit**	**Strategy**
Folgers	1[b,c]	**92**	3	4	$1,093	EDLP
Maxwell House	**47**	0	3	**50**	$804	Wide Pulsing
Chock Full O' Nuts	2	**91**	3	4	$527	EDLP
Pattern 2, 16 runs	**1**	**2**	**3**	**4**	**Average Profit**	**Strategy**
Folgers	1	**94**	2	4	$1,092	EDLP
Maxwell House	**47**	1	3	**48**	$804	Wide Pulsing
Chock Full O' Nuts	1	**91**	3	4	$527	EDLP
Pattern 3, 1 run[d]	**1**	**2**	**3**	**4**	**Average Profit**	**Strategy**
Folgers	2	**92**	0	6	$1,045	EDLP
Maxwell House	**46**	0	4	**50**	$830	Wide Pulsing
Chock Full O' Nuts	**48**	**44**	4	4	$580	Promote to the Max

a: patterns of competition are computed during the 100th generation from all combinations of 25 agents playing 52-week games and from 50 Monte Carlo simulations.

b: row percentages, key numbers in bold.

c: shaded areas identify the actions constrained by store policy.

d: best performing of remaining patterns.

population case, 11%, 6% and 6%. We believe that this increased heterogeneity is a result of the absence of the inbreeding we observe in a common population, where brands exchange information at the genetic level through the recombinant processes of the GA—“incest,” as Tony Curzon Price has put it (personal communication).

Experiment 2 — four strategic players

The results of expanding from three to four players—by adding *Hills Bros*.—are summarized in Figure 2, along with those from Experiment 1 for comparison. We again use four actions but now with four distinct populations. Whilst *Hill Bros.* is a niche player, and has smaller profits than other brands, its inclusion produces complex changes in the behavior and performance of the major brands. *Folgers* switches from EDLP to an aggressive strategy alternating between its EDLP price and a featured price promotion. *Maxwell House* pulses between two high shelf prices instead of using a featured-price promotion as before. *Chock Full O’Nuts* switches from EDLP to pulsing between a high shelf price and a featured-price promotion. *Hills Bros.* pursues a similar strategy to *Maxwell House.*

Figure 2 about here

Not only are the changes in agent behavior significant, so too are the shifts in profitability. As Figure 2 shows, in comparison to the three-population case, the introduction of *Hills Bros.* reduces the profits of *Folgers* and *Maxwell House* (by 13% and 8% respectively) but increases the profits of *Chock Full O’Nuts* (by 56%).

One reason for the falls in profitability experienced by *Folgers* and *Maxwell House* are that the agents for *Hills Bros.* are frequent price promoters and take up promotional opportunities that would otherwise have gone to one of these major brands. With the store policies set the way they are, there is only one opportunity for a featured price promotion each week. The greater the number of brands competing for these opportunities, the fewer the opportunities any one brand is going to have. Moreover, as these promotions usually generate significant profits, *ceteris paribus*, the overall profitability of the brand is reduced. Another reason is that *Hills Bros.* promotions themselves take sales from the other brands.

Figure 2 Four-Action Experiments

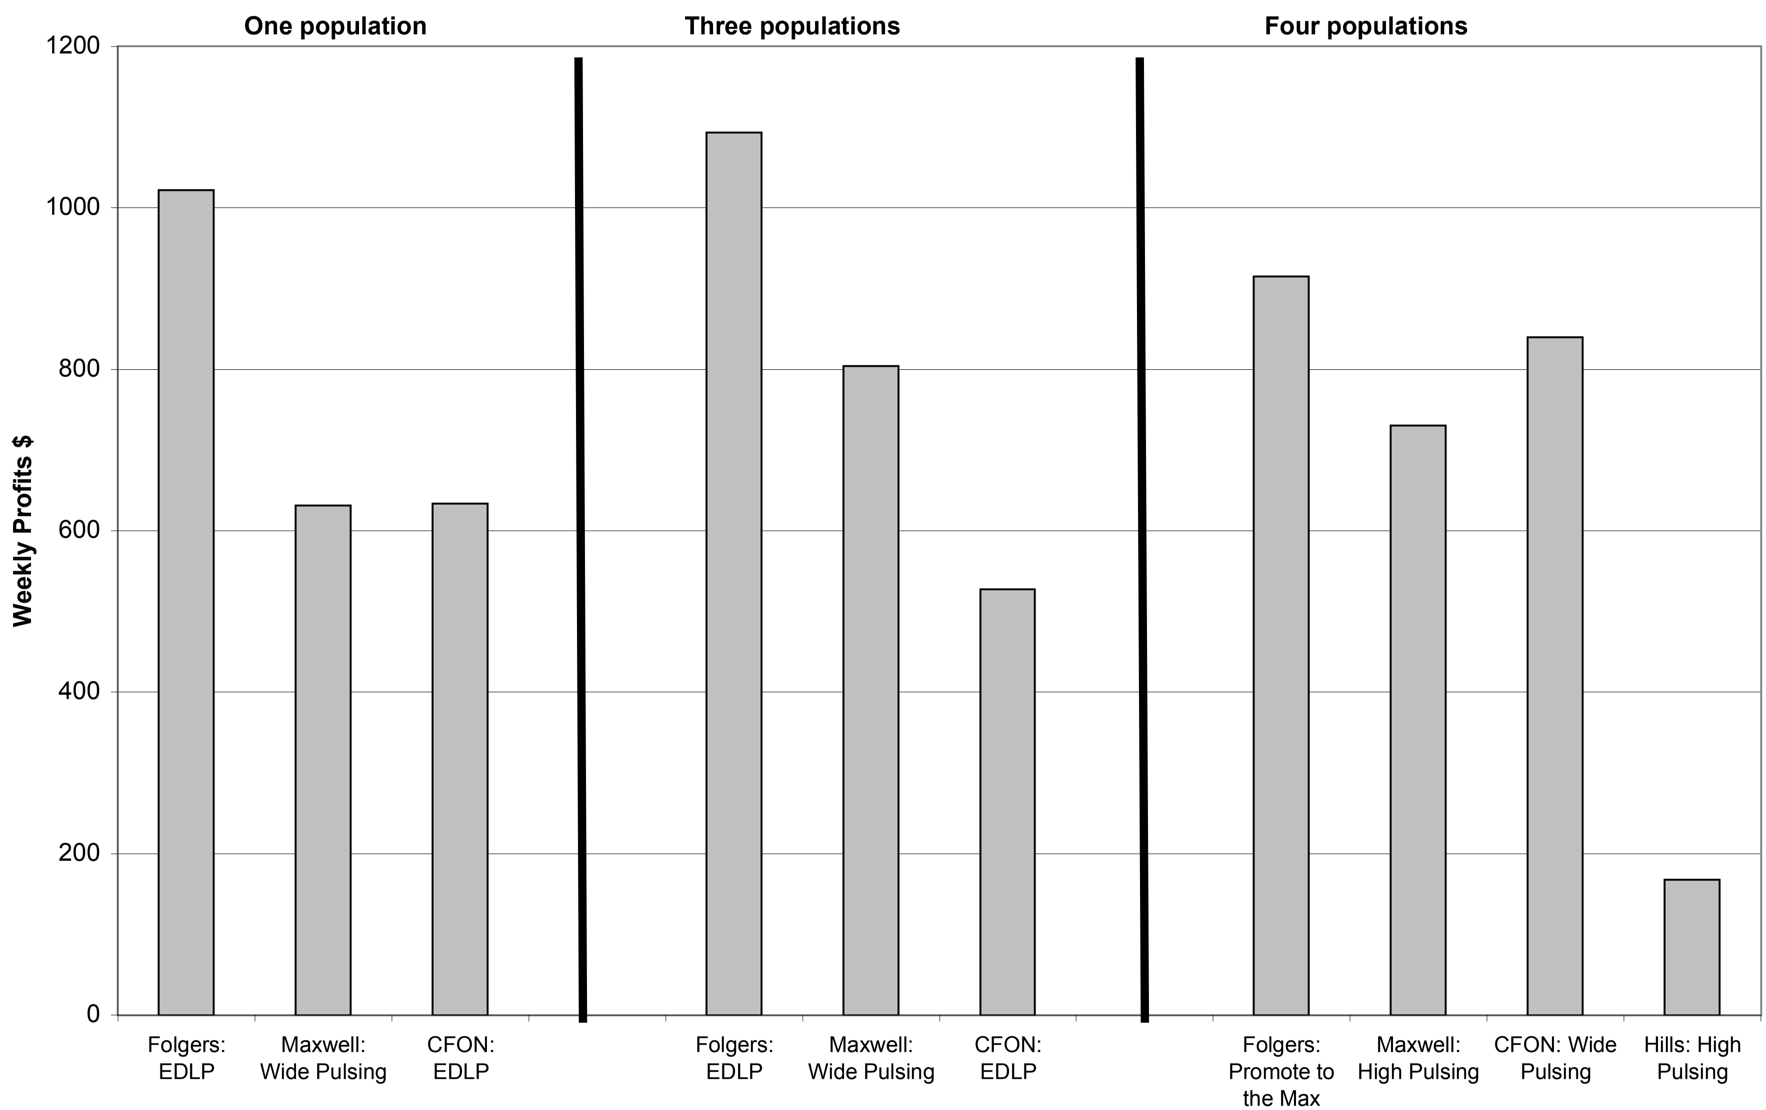

But how do we explain the increase in profitability experienced by *Chock Full O'Nuts*? Examination of the parameters of the market response model reveals that the actions of *Hills Bros.* have less impact on the sales of *Chock Full O'Nuts* than the other brands (Cooper and Nakanishi 1988, Chapter 5). The introduction of *Hills Bros.* as the fourth player thus benefits *Chock Full O'Nuts*—because each *Hills Bros.* promotion damages the other brands but has less effect on *Chock Full O'Nuts*.

Overall, the introduction of a fourth player had a greater impact than we anticipated. One of the strengths of our approach is the use of a detailed, realistic and empirically grounded model of consumer response. The realities are that this response, and the brand competition that occurs as a consequence of it, are complex phenomena. Like managers, the agents vigorously compete with the tools they are given. In the four-action, four-competitor setting this provided an advantage for the agents representing *Chock Full O'Nuts*. But, as we shall see in Experiment 3, when the agents are given eight actions, the competitors of *Chock Full O'Nuts* change their behavior to better overcome the inclusion of *Hills Bros.*

Experiment 3 — eight possible actions per player

Increasing the number of possible actions provides the agents with more choices of how to compete—both in terms of the number of available actions and because we have extended the range of possible prices in comparison to the four-action simulations.

Early in the co-evolution, the strings tend to use each of the eight actions with a roughly equal frequency. However, by the 100th generation the agents are using fewer than eight actions. In fact the one, two, three or at most four actions for each brand account for between 66% and 85% of all actions. The agents have learned the actions that are most profitable for them, given the behavior of their rivals.

These learned actions, however, vary significantly between brands and somewhat across the two patterns that our Monte Carlo runs fall into. Thus in the three population case the agents for *Chock Full O' Nuts* use a similar strategy across both competitive patterns—in essence 'pulsing' between a high price and a featured low price. *Folgers* also pulses, but in Pattern 1 (27 of 50 simulations) the agents focus on four actions and in Pattern 2 (14 of 50 simulations) only three. *Maxwell House* also uses a similar strategy across both patterns, namely two variants of an everyday-low-price strategy. Overall, these results indicate the importance of store policy. This can be seen in the various ways the agents adapt their behavior to these constraints. For example, *Maxwell House* chooses the lowest unconstrained action for its everyday low price, whilst the other two brands only select certain of the constrained actions for their promotional pulses. Figure 3 summarizes the eight-action results for both three and four players. Two aspects of these results are worth commenting on.

Figure 3 about here

First, whereas with four possible actions we typically only observe one pattern of competitive interaction across the 50 Monte Carlo runs, with eight actions we observe two patterns. Whilst the differences between the two are more a matter of degree than of kind, there is no doubt these results show more heterogeneity than do the four-action experiments. In the case of three players, as discussed above, the differences primarily concern whether the agents for *Folgers* focus on three or four actions. For the case of four players, in 9 of 50 runs we observe the High Entropy strategy across all brands. This strategy has its actions fairly evenly distributed between four to six of the eight possible actions. In 30 of 50 runs we observe all brands placing a greater focus on three middle to high-priced actions—which we term High Pulsing. The fact that we observe two patterns in both the three- and four-player, eight-action experiments indicates that different patterns of competitive interaction can appear over long periods of co-evolution. Since we did not observe this in the equivalent four-action experiments, we conclude that different paths of co-evolution are more likely as we increase the number of actions available to our agents.

Second, again it can be seen that the introduction of *Hills Bros.* produces significant changes in the behavior of the major brands. Compared with the three-player, eight-action experiments, the profits of all the major brands fall, although it is noticeable that *Chock Full O'Nuts* and *Hills Bros.* perform better in the High Entropy strategy. In the High Pulsing condition the other brands are able to offset these advantages by clearer focus on certain actions. In the four-action experiments they were not able to overcome these advantages. Thus a broader palette of actions may enable more effective competition.

Figure 3 Eight-Action Experiments

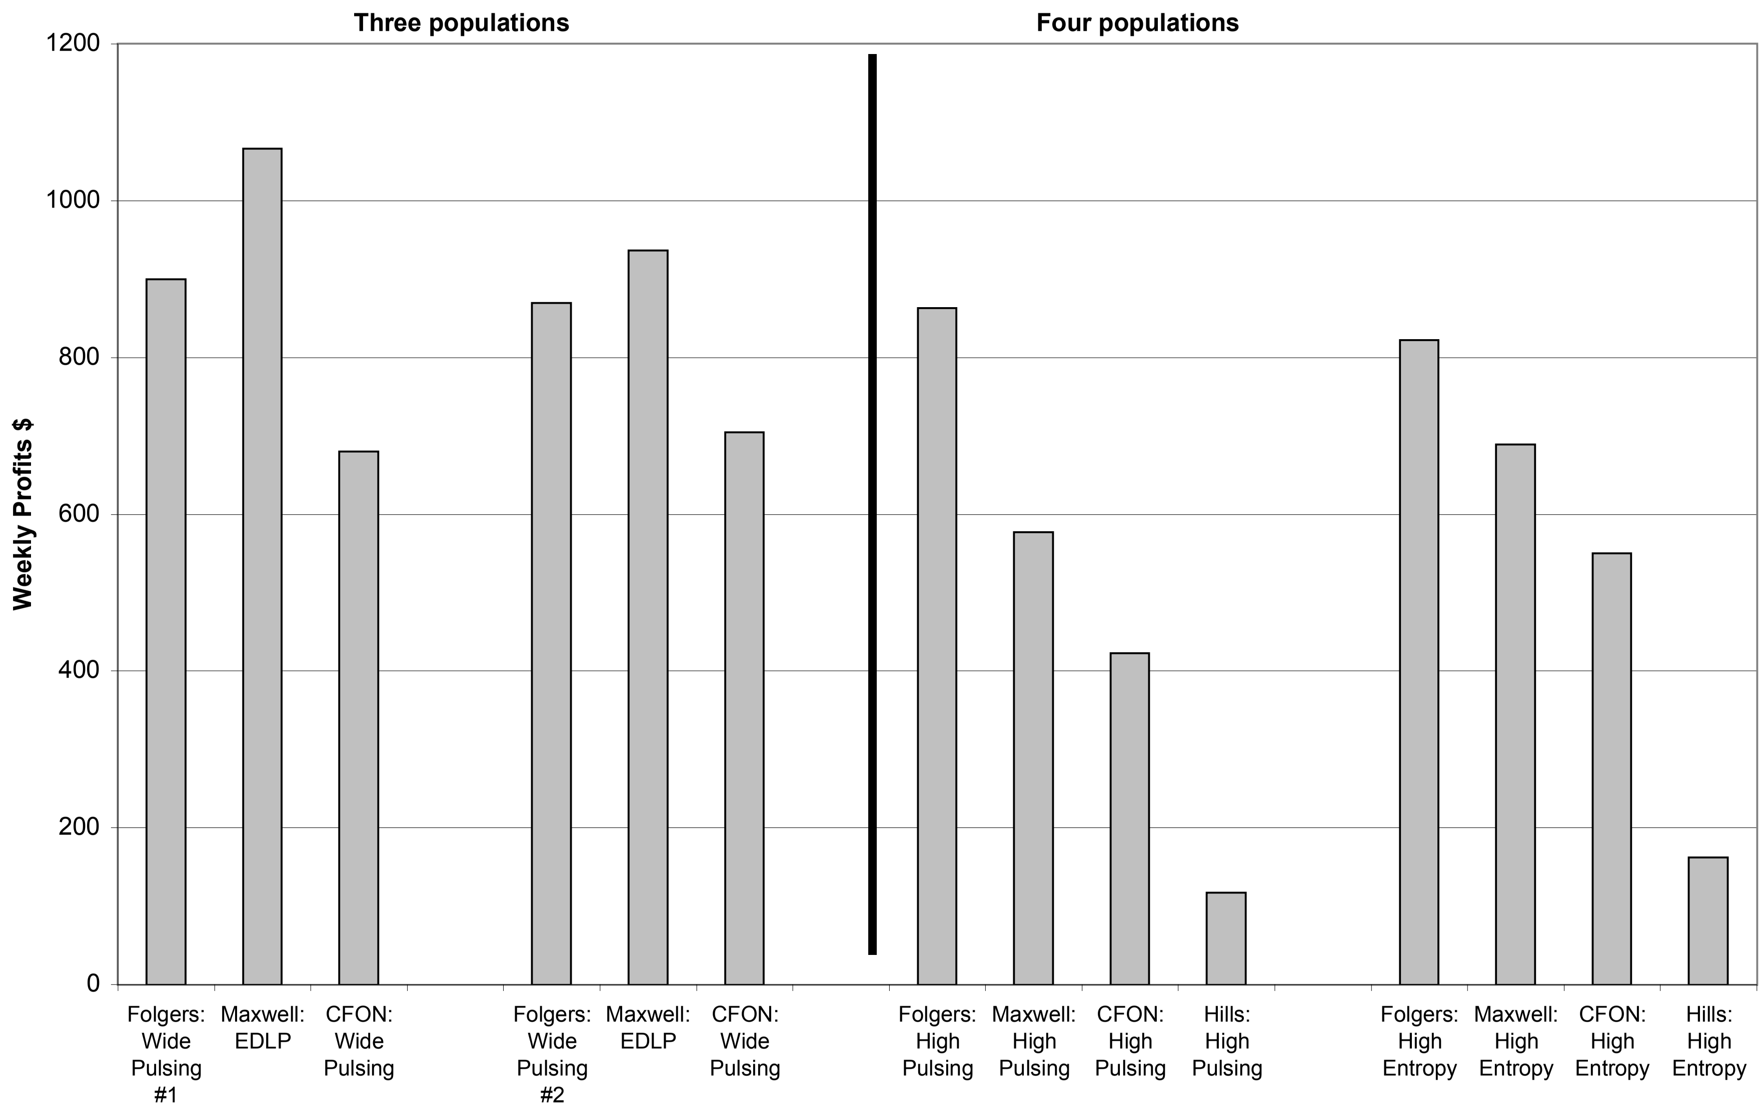

Experiment 4 — co-evolution: sophisticates against primitives

We want to demonstrate that the natural selection of the GA is improving the performance of the artificial agents, which start off as filtered (and hence legal) but otherwise random strings. The improvement of the mean score of the population as shown in Figure 1 is one measure, the good performance of the best agents against the historical actions of managers is another. In this section, we consider yet another measure: pitting the best-performing string after 100 generations (the 'sophisticate' agent) against all combinations of the other two brands after 8 generations (the 'primitive' agents). What should we expect? Since the sophisticates have had many more generations to learn and adapt than have the primitives, we should expect them to score better against the primitives than against the sophisticates. The results are given in Table 3.

Table 3 about here

The Holyfield-Tyson effect. We should expect positive diagonal elements in the table: that is, the sophisticate's performance would be significantly better than that of the replaced primitive string. But despite our expectation, only with one brand (*Chock Full O'Nuts*) does this occur; the other two brands show falls in the average weekly profits of between 3% and 5% from the level of the replaced primitive.

From this table we are left to conclude that, at least for two of our three brands, the sophisticated agents do not compete effectively with primitive agents. By analogy to a famous ear-biting incident, Bernhard Borges has called this the Holyfield-Tyson effect (personal communication). Without a referee to call fouls, the primitive agents turn the competition into a street brawl. One explanation for this effect might be genetic drift.

Genetic Drift. Genetic drift is the change in the gene pool of a small population that takes place strictly by chance. Genetic drift can result in genetic traits (genes, or patterns on our bit strings) being lost from a population or becoming widespread without respect to the survival value of the genes involved. As a random effect, genetic drift can occur only in small, isolated populations in which the gene pool is small enough that chance events can change its makeup.

Table 3. Mean changes in average weekly profits with best sophisticate competing against the best primitivesa

Best Sophisticate	Change in *Folgers*	Change in *Maxwell House*	Change in *Chock Full O'Nuts*
Folgers	-15.01	41.42	42.03
Maxwell House	2.03	-20.04	37.77
Chock Full O'Nuts	13.93	-28.99	82.34

a: this table is based on the following 8 step procedure.

1. After eight generations, identify the best individual strings from each of the three populations.
2. Play these three against each other for a 50-week repeated game, and note their average weekly profits.
3. Allow the three populations to continue co-evolving via the GA.
4. After 100 generations, identify the best strings from the three populations, play them against each other as before, and note their average weekly profits.
5. Replace the best *Folgers* string after 8 generations by the best *Folgers* string after 100 generations (e.g. replace the best primitive string by the best sophisticate).
6. Play all combinations of three strategic brands, and consider string-by-string the change in average weekly profits with the sophisticated player and without the sophisticated player in one brand.
7. Repeat steps 5 and 6 for the remaining two brands.
8. Because of the stochastic nature of the simulation, repeat steps 1 through 7 fifty times and compute the averages shown in the table.

In larger populations, any specific gene is carried by so many individual strings that it is likely to be transmitted to the next generation, unless it scores badly.

The magnitude of gene frequency changes due to genetic drift is inversely related to the effective size of the population (the number of individuals selected to produce offspring), since only parent strings transmit their genes to the following generation. In our case we have a small population to begin with, only 25 individuals, and the number of parents in each generation is smaller still, so that genetic drift over 100 generations may become significant. It is possible that we have lost genes that led to higher weekly profits when the primitives were competing from the population, or, in the vernacular, that the sophisticates have become "flabby" in competition against similar sophisticates.5 This could explain the poor showing that the *Folgers* and *Maxwell House* sophisticates make against the primitives; the *Chock Full O Nuts* sophisticates consistently do better against primitives than against sophisticates.

In order to test the conjecture that the sophisticates were not doing as well against the primitives because of genetic drift, we decided to increase the size of each population from 25 to 250. The tenfold increase in the population sizes means that simulations take much longer, since each individual now has to compete against 250^2 combinations (instead of 25^2) and there are ten times more individuals to test per generation. This thousand-fold increase in the number of three-way interactions per generation also means that convergence is much slower. It is not just that each of the three populations takes longer to converge (as it would by itself) but that the slowness of the opponents to converge means that any convergence that does first occur is likely to be premature. With the three populations converging roughly at the same rate, *Folgers*, for instance, may find a small set of strings whose patterns mean good profits, against the current populations of *Maxwell House* and *Chock Full O'Nuts*. But when further convergence in, say, *Maxwell House* occurs, the *Folgers* strings may be further from optimal, which means that convergence may reverse, until the *Folgers* population adapts to the new topology in strategy space, and evolutionary convergence continues. We can think of this as a spiraling towards a node: closer but then further away, and then closer, and so on. We have observed this spiraling in the lengthy simulations we have run on populations of 250 (which take months rather than hours to complete). We performed a single run only and we obtained the same results as before—namely sophisticates perform poorly against primitives.

These large-population simulations would appear to rule out genetic drift as an explanation for the Holyfield-Tyson effect, but we must qualify this null result in two respects. First, the cycles of convergence are much longer in large populations than in the smaller populations typical in GA applications. It is possible that, were we to continue the simulations for more generations, we might obtain sophisticates capable of holding their own with primitives. Second, the immense number of computations involved has meant that we have only been able to investigate genetic drift for a model with three populations and four possible actions. As discussed above, a larger palette of actions appears to result in more capable agents. It is possible that a sophisticate with eight possible actions might be capable of competing with a primitive with eight actions. Our intuition is that neither of these potential explanations will prove valid—but we need to note that we have not tested them.

Experiment 5 — managerial learning

Thus far we have looked at how agents learn in competitive simulations, but our methodology also allows us to look at how managers may have learned over decades of actions in the coffee market. In Experiment 3 we used eight actions that were developed from analysis of the historical data, that is from managers' actions. These eight actions are thus highly learned ones and exogenous to the agents that we evolve in our simulations. All the agents detailed above are similarly restricted to a repertoire of actions we specify for them. The question arises of how these agents might perform with a different repertoire of actions—one developed without reference to the historical actions of managers. These we developed from a random experimental design where price is stepped in ten-cent increments between $1.60 and $2.80 and feature and display can take on the value of either 0 or 100.

In Table 4 we show the patterns of competitive interaction that are observed across the 50 Monte Carlo simulations we ran with these eight randomly chosen actions. In this case there are two patterns that account for 31 out of 50 runs. We can note two striking facts about Table 4. First, the profit levels are much higher than with the earlier, learned actions. Second, these profit levels are achieved because the agents are very sparing in their use of featured, low-price promotions and maintain high prices throughout most interactions. This is certainly true for *Chock Full O'Nuts,* whose agents stay at high prices most of the time. The agents for *Folgers* 'pulse' in both patterns but also maintain the highest price more frequently than in earlier results. Similarly, *Maxwell House* has a definite high-to-low pulsing strategy but also stays at the highest price more frequently. In essence, the level of competition is much lower with these random actions. And on first glance it would appear as though managers have 'learned' to become overly competitive in comparison to these agents.

Table 4 about here

We speculate, however, that what these results show is that store competition is missing from our response model and, we would add, from much of the market-response modeling literature. In the market, competition from other chains may prevent such high prices. This competition is likely to affect store policies, and the impact of these policies may be incorporated into the historical actions of the managers. Thus, when we extrapolate beyond the bounds of our earlier simulations, as we have done in Experiment 5, we see behavior that may not be sustainable in the market. Or indeed, may not be observed in the market—since the mere threat of competition may also alter the beliefs of store managers.

Agents and managers

Our final section of results compares the agents we have bred with the historical actions of brand managers. We did this in two ways. First, we took the best agent from each of the distinct populations and played these against the actions of the managers of their competitors.[6] The profits achieved by these agents could then be compared with the historical profits achieved by managers of the same brand. This procedure is similar to the one we used in our earlier paper, but, whereas previously we used 25 agents from one common population, here we use the best agent from each of 50 brand-specific populations of 25 agents each. Whilst we recognize that this test is unfair—in that the actions of the competing managers are frozen—it does give us another benchmark for agent development. If our agents cannot compete with these managers we might well question their value. Second, we know that whilst agents

Table 4: Patterns of competition among evolved agents—three distinct populations and eight random actions

	Actions **Low price** **High Price**								
Pattern 1, 18 runsa	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**Profit**
Folgers	**24b, c**	14	1	1	0	8	2	**47**	$2,009
Maxwell House	14	**24**	2	1	1	8	3	**45**	$2,801
Chock Full O'Nuts	4	1	1	0	2	6	**29**	**52**	$1,151
Pattern 2, 13 runs	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**Profit**
Folgers	5	**34**	0	4	0	4	3	**47**	$1,916
Maxwell House	**41**	2	0	2	0	3	2	**48**	$3,244
Chock Full O'Nuts	4	1	1	1	1	2	3	**85**	$1,192

a: patterns of competition are computed during the 100th generation from all combinations of 25 agents playing 52-week games and from 50 Monte Carlo simulations.

b: row percentages, key numbers in bold.

c: shaded areas identify the actions constrained by store policy.

competing against themselves produce higher profits than managers they do so with different patterns of promotion. It is therefore interesting to contrast the behavior of the agents with that of managers.

Agents against history. On average, for the four-action simulations, these 50 best agents outperform *Maxwell House* managers by 12% compared with –15% in our previous paper. Moving to a distinct population for *Maxwell House* produces improved agent profitability. Similar tests for the other two brands produce average figures of 121% for *Folgers* (compared with 49% in the previous paper) and 55% for *Chock Full O' Nuts* (compared with 53%). Whereas *Maxwell House* was the main beneficiary of distinct populations when competing against other agents, it appears that *Folgers* is the main beneficiary when competing against managers. For the eight-action simulations, the agents for *Folgers* outperformed the historical actions of its managers by 156%, *Maxwell House* by 32% and *Chock Full O'Nuts* by 42%. For *Folgers* and *Maxwell House* these results are better than those from the four-action experiments, for *Chock Full O'Nuts* they are slightly worse.

The Frankenstein effect. A more interesting finding from these tests against the historical actions of managers is that the agents here exhibit more heterogeneous behavior than when they compete against other agents. For example, in the four-population, eight-action experiment, we observe two patterns of competitive interaction when agents compete against agents. When agents compete against managers, however, we observe two patterns for *Folgers*, but four for *Maxwell House*, *Chock Full O'Nuts* and *Hills Bros*. Similar results hold for several of our other experiments. The historical actions likely trigger more diverse behavior from the agents because the historical actions themselves contain more variability of actions than we see from the agents after 100 generations of co-evolution. We find it interesting and important that the artificial agents can handle these changed inputs, perform well, and do so with strategies different from those they used when competing with each other. We call this the Frankenstein effect—agents that only showed a few behaviors in the 'laboratory' show more in real environments.

Comparing the behavior of agents with that of managers. In Table 5 we show three sets of action frequencies for each of our four brands. The first set is the historically observed frequencies for these brands, the second is from our four player, eight action simulations and the third is from pitting the best agents from these simulations against the historical actions of competing managers.

Table 5 about here

Examining these frequencies, particularly the emboldened figures, it is evident that the agents use featured price promotions more than managers and they use the highest price more (especially when competing against the historical actions of managers). The net effect of their behavior is to reduce the average price of coffee (by 5% when competing against each other and 4% when competing against managers) and to increase the average amount of coffee sold by the chain each week (by 40% and 9% respectively). Whilst these numbers represent significant changes to the market they are not unreasonable. Historically, prices and amounts sold have varied by more this. For example, the maximum weekly amount sold by the chain was 171% of the annual average. Moreover, even for agents competing against themselves—the most vigorous price competition we observe—the implied amount sold only represents 46% of the average amount of coffee sold in the regional market in one week.

The key question is thus not whether customers can consume this amount of coffee—they clearly can—but whether agent behavior can be sustained in the long run against inter-chain competition. That is, might the other chains in the area adjust their policies to address the potential loss of business that this behavior represents? While this is not a straightforward decision for them, our intuition is that they would—which another reason why we think that inter-chain competition is an important area for future research.

Discussion

In summarizing our findings and conclusions, we have chosen to touch on four main areas: co-evolution, the specification of agents, modeling more brands, and managerial behavior. After discussing these, we will briefly outline our future research plans.

Table 5. Comparisons of the behaviors of agents and managers

	Actions							
	Low price							**High price**
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Folgers								
Historicala	1	1	1	5	2	3	**87**	1
Agents competingb	6	5	8	**14**	3	**39**	9	**16**
Agents against historyc	**11**	7	**12**	7	9	7	7	**40**
Maxwell House								
Historical	1	7	3	**11**	6	**61**	**11**	1
Agents competing	5	6	6	**16**	6	**39**	9	**13**
Agents against history	8	**13**	**10**	9	**10**	6	5	**40**
CFON								
Historical	1	**13**	3	2	1	7	**66**	7
Agents competing	8	6	6	**15**	4	**39**	**10**	**12**
Agents against history	8	7	8	**13**	**10**	**12**	7	**35**
Hills Bros.								
Historical	1	2	8	3	9	**69**	8	1
Agents competing	6	8	5	**12**	5	**39**	9	**16**
Agents against history	9	8	7	**12**	9	6	5	**44**

a: from historical data, row percentages, key numbers in bold.

b: from the most common pattern observed in the four population, eight action simulations

c: from playing 50 best agents against the historical actions of their competitors.

Co-evolution

The methodological contribution we make in this paper is the rewriting of the GA to co-evolve multiple populations of agents. This is clearly a more valid approach for our application. Each brand evokes a different response from consumers, and faces different cost structures, and these facts should be reflected in the agent's mapping from perceptions to actions. While a single-population GA can produce a solution, it is a 'one size fits all' solution, whereas a multiple-population GA allows a distinct population of solutions for each brand. Moreover, the single-population GA pools genetic information from the competing brands to an undesirable extent. The empirical confirmation of these arguments is that the agents evolved from our three- and four-population simulations outperform the common-population agents of our earlier paper.

But co-evolution incurs some additional costs, because it requires more computations to reach convergence. This is partially because of the computations involved in keeping track of the multiple populations, but extra computing also occurs because the multiple populations do not evolve in step. Bit strings for one population can converge only to find that their competitive environment (the other populations) has changed. They must then evolve again to catch the moving target. Overall convergence is thus approached in a 'spiraling' manner and can take longer than with a single population. Indeed, it is only because we were able to improve the GA code, and because computer hardware has improved since our earlier work, that we are able to conduct these multiple-population simulations.

We also have an unresolved problem with the results of co-evolution—our finding that 'sophisticated' agents do not compete well with 'primitive' agents. In one sense this is not a problem. Agents perform well against other agents at the same level of sophistication. Eventually, however, we would like to breed agents that are robust to changes in competition For example, the introduction of a new 'primitive' player into a market. We need bit strings that retain the capacity to compete with agents other than those they currently face.

Our initial supposition that the loss of this capacity was due to genetic drift in a small population is not borne out by our results. We have increased the population size tenfold but still obtain similar results to before. Of course, even this increased population is not a large by the standards of the natural world, but our intuition is that increasing this population yet further will not solve the problem. We may need to look at other methods for retaining information in the bit strings, such as a diploid mechanism. The diploid genotype (Goldberg & Smith 1987) (which allows the emergence of "recessive genes", such as that for blue-eyed human beings) might allow once and future genes of value to persist, even in an ever sophisticated population. Mitchell (1996, pp.21-27) discusses work by Hillis, in which he used a diploid representation: chromosomes in pairs, rather then the single chromosomes or haploid representation that we use here.7

Specification of agents

While we still have an unresolved issue with our basic genetic mechanism, the results presented here do allow more definite conclusions to be made about the form and structure of our agents. We have learned that our agents do not need as many actions as perhaps we had thought or as perhaps managers use.

Our earlier work used four possible actions. At the time we thought that this was a reasonable number with which to start our work. It is also true that the state of development of our methods and the available computing hardware had some impact on this choice—it would have been very time consuming to investigate a larger number. Here advances in our methods (for example, 'filtering'), and in the computing power available to us, have allowed us to investigate eight possible actions. But we find that the agents end up using far fewer than eight actions. In fact they most often end up using two or three with high frequency and the rest with low frequencies. And we only see incremental rather than dramatic increases in the profitability of these agents. This raises the interesting question of whether we need to equip our agents with as many actions as we do in this paper.

Of course, we need to be careful in this assessment—a low-frequency action may still be important to the overall performance of the agent. But it does suggest that eight is a reasonable upper bound and we need not equip our agents with more possibilities.

We have conceptualized the structure of our agents in three components: 1) a mechanism for perceiving the previous state of the market, 2) a mechanism for mapping from this state to future actions, and 3) a mechanism for determining an appropriate set of actions. Our earlier paper addressed simply the mapping mechanism using exogenously determined perceptions and actions (which were in fact equivalent). This paper goes some way to addressing the third component by suggesting an upper bound of eight possible actions. Our intuition is that the most productive areas for future research are perceptions, and methods for endogenizing the development of perceptions and actions.

Modeling more brands

In our previous paper we modeled competition between the three major brands in this market, brands whose combined market shares account for some 68% of the market in question. Here we introduced agents representing a fourth, niche brand whose market share is only 4% and whose profitability is much lower than the major brands. Yet the introduction of this niche brand changes the market in significant, complex and asymmetric ways. Not only does it impact on profits but it also changes the behavior of the major brands. To some extent this occurs simply because the fourth brand takes up some of the fixed number of opportunities for major promotions—thereby forcing the agents for other brands to reallocate their actions. But it also occurs because the fourth player has different competitive impacts on each of the three major brands. This disturbs the competitive balance between these brands, and to compensate for this the behavior of their agents also changes. Which raises the question of how many brands we need to include in our analysis. Before this finding, we would have answered “just the major brands”, but now we are not as certain. Conceivably, researchers need to model all the brands in a market to fully understand the nature of competition in that market. Fortunately, the continued increases in computing power make it possible to contemplate such simulations in the not too distant future.

Managerial behavior

In most of our work we have exogenously determined the actions of the agents by analysis of the historical actions of managers. Since we also equivalence perceptions to actions, these too are based on the historical data. In that sense our agents have ‘learned’ from, or at least been informed by, what managers think are appropriate competitive actions. The agents are thus performing well because the GA evolves effective mappings between perceptions and actions. And we have some evidence that these mappings are at least as effective as those implicit in the managers’ behavior.

In our last experiment we broke from this path—by randomly drawing a set of possible actions for the agents. Our prior hypothesis was that agents with these random actions would not perform as well as those with exogenously determined actions; in other words, that the managerial learning incorporated in the exogenous actions had value. We were therefore disconcerted to find that the agents with random actions extract much higher profits from this market. Is this because we have been lucky with our random draw? Or is there some miss-specification in our model?

Our agents, faced with a menu of random actions, make most of their profits at prices above their shelf prices—the highest prices used by managers. Such outcomes are what one might have expected before the data used in our earlier paper had been examined: that the highest profits earned would occur at collusive, high prices. But those data show that the greatest profits earned by the historical brand managers occur, not at high, above-shelf-price levels, but at low, promotional prices. This suggests that one reason for the high profits of agents with a menu of random actions is that the demand in our models for coffee at high prices is not constrained by competition from other supermarket chains. We have not yet modeled the effects of this competition on the demands faced by our agents. Once we have done so—an admittedly difficult exercise—we may be better able to assess the behavior of managers in comparison with agents. All in all, Experiment 5, together with our subsequent comparison of agent and manager behavior, reinforces our belief that there is much to be gained by jointly examining the competition between retailers and manufacturers.

Future research

A conclusion we draw from all the above is that it would be better if our agents endogenously determined their own perceptions and actions. Then we could compare the performance and behavior of these agents with that of managers without the confounding of the two which comes from our current selection of actions and perceptions. But this is a non-trivial task. One of us has done some thinking on partitioning perceptions (Marks 1998) but we lack a specific algorithm to implement. We have not, as yet, done much thinking on what would be a basis for selecting a set of actions. Perhaps parts of the bit string could represent a price or promotional level. Similarly, perceptions might be represented as partitions—the coordinates of which are encoded in the bit string. The GA would then exert selective pressure on these bits, resulting in the selection of partitions and actions that maximize profits.

The other area we wish to work on is that of the interaction between store policy and brand management. It is clear from our results that store policy has a major impact on brand competition. We need to consider joint models whereby store managers seek to impose policies that maximize their profits whilst brand mangers seek to maximize profits within these policies. Presumably store managers also consider competition from other store chains and their beliefs about this competition are reflected in their policies. Since brand managers also seek to influence store policy by various incentives, these will be complex models.

To conclude, we note that we believe the strength of our overall approach derives from the use of an empirically grounded fitness function to optimize our agents. A market-response model based on a real market; such as we use here, provides a strong test of the validity of the artificial agents. It also focuses our research endeavors on realistic specifications for these agents. While empirical grounding is not yet common in applications of genetic techniques, we suggest that such grounding has much to commend it.

1 The response model comprises a market-share sub-model, a category-volume sub-model, and a profitability sub-model. It uses over 180 parameters and explains the historical volume and share levels very well. Brand price elasticities vary from –0.5 to –4.7 and costs between $1.18 and $1.39 per pound. Demand saturation is modeled by pro-rating a moving average of the amounts sold against a saturation limit (the average amount of coffee sold in the region).

2 Eight actions requires 3 bits per action; 8 actions, 4 players, and 1-week memory implies 84 = 4,096 possible states; phantom memory is 8 x 3 = 24 bits. So 3 x 4,096+24 = 12,312 bits.

3 The particular store chain that we model does not allow brands to feature two weeks in a row, and it only allows one brand per week to feature. In the GA we impose these constraints by penalizing the profits of actions that would violate them.

4 "Nothing is absolutely predictable about the direction of co-evolution. How an interaction co-evolves depends not only on the current genetic makeup of the species involved but also on new mutations that arise, the population characteristics of each species, and the community context in which the interaction takes place." *Encyclopaedia Britannica*, CD97.

5 An analogy in game theory is the issue of how a strategy will perform against an irrational strategy that moves off-equilibrium.

6 E.g. an agent determines the actions for *Folgers* whilst all other actions are historical.

7 Since there are pairs of chromosomes, at each position on the first chromosome the string segment is compared with the corresponding string segment on the second chromosome. If they encode the same action, then that action is the contingent action for the corresponding state; if they encode different actions, then one of the actions will be dominant, and the other recessive. The coding for the recessive action may survive the recombination of the chromosomes during the GA's evolution, and eventually may prove valuable.

References

Axelrod, R., “The evolution of strategies in the iterated Prisoner's Dilemma”. In: L. Davis (ed.) *Genetic Algorithms and Simulated Annealing*, L. Davis (ed.), Pitman, London, 1987, 32-41.

Cooper, L.G., and M. Nakanishi, *Market Share Analysis: Evaluating Competitive Marketing Effectiveness*, Kluwer, Boston, 1988.

Everitt, B.S., *Cluster Analysis*, London, Heineman, 1980.

Goldberg, D.E., & R.E. Smith, “Nonstationary function optimization using genetic algorithms with dominance and diploidy”. In: Grefenstette J.J. (ed.) *Genetic Algorithms and their Applications*, Proceedings of the 2nd. International Conference on Genetic Algorithms. Lawrence Erlbaum, Hillsdale, N.J., 1987.

Marks, R.E., “Breeding Optimal Strategies: Optimal Behavior for Oligopolists,” *Journal of Evolutionary Economics*, 2 (1992), 17-38.

Marks, R.E., “Evolved perception and behavior in oligopolies,” *Journal of Economic Dynamics and Control*, 22, 8-9 (1998), 1209-1233, forthcoming.

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

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